



HAL
open science

Heat diffusion distance processes: a statistically founded method to analyze graph data sets

Etienne Lasalle

► **To cite this version:**

Etienne Lasalle. Heat diffusion distance processes: a statistically founded method to analyze graph data sets. *Journal of Applied and Computational Topology*, 2023, 10.1007/s41468-023-00125-w . hal-03366848

HAL Id: hal-03366848

<https://hal.science/hal-03366848>

Submitted on 3 Jul 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

HEAT DIFFUSION DISTANCE PROCESSES : A STATISTICALLY FOUNDED METHOD TO ANALYZE GRAPH DATA SETS

Etienne Lasalle,
Laboratoire de Mathématiques d'Orsay,
CNRS, Université Paris-Saclay,
Orsay, France
etienne.lasalle@universite-paris-saclay.fr

ABSTRACT

We propose two multiscale comparisons of graphs using heat diffusion, allowing to compare graphs without node correspondence or even with different sizes. These multiscale comparisons lead to the definition of Lipschitz-continuous empirical processes indexed by a real parameter. The statistical properties of empirical means of such processes are studied in the general case. Under mild assumptions, we prove a functional central limit theorem, as well as a Gaussian approximation with a rate depending only on the sample size. Once applied to our processes, these results allow to analyze data sets of pairs of graphs. We design consistent confidence bands around empirical means and consistent two-sample tests, using bootstrap methods. Their performances are evaluated by simulations on synthetic data sets.

1 Introduction

Considering the current growth of available data and the modeling power of networks, methods to analyze graph-structured data have gained interest over the last few decades. Particular attention has been devoted to designing notions of distance between graphs. The design of these notions is highly constrained by the working framework. In particular, different types of information can be used depending on whether the graphs are directed or undirected, weighted or unweighted, have the same size or not. Another key factor to define distances is whether a node correspondence (NC) is known or not. In the case of known NC, we can consider the graphs to be defined on the same vertex set and comparisons can be made at the edge scale. In this context, people have applied various metrics to compare adjacency matrices, Laplacian matrices, heat kernels [20] and other matrices whose entries represent quantities associated to pairs of nodes [26]. On the other hand, when no NC is available, graphs are often compared at a mesoscopic or macroscopic level using structural summaries. People have used global statistics on graphs like degree distributions, network diameters, or clustering coefficients [36]. Another well studied approach is to consider graphlets [36], *i.e.*, small given subgraphs that are counted in graphs. Then, various methods have been developed to compare the graphlet counts [35, 49, 1, 14]. Some work has also been pursued to exploit the structural information carried by spectra of operators [48, 18]. To measure distances (or equivalently similarities) between graphs, one can also rely on graph kernels, hence benefiting from the general kernel methods to solve statistical problems. These various kernels can be based on neighborhood information, subgraphs structures, random walk properties, shortest paths. Unfortunately, these kernels often require additional information like node or edge labels, attributes, which are not always available. Moreover, not all graph kernels are able to handle weighted graphs. For more details on graph kernels, see the survey of [27].

Another powerful way to encode structural information about graphs is to use diffusion processes, like heat diffusion. When working with weighted graphs, one can interpret weights as the thermal conductivity of edges, meaning that heat diffuses faster along edges with higher weights. Note that unweighted graphs can always be seen as weighted graphs with weights in $\{0, 1\}$. Given initial conditions, the way heat diffuses can be used to characterize and compare graphs [10, 9, 20, 45]. This approach is appealing as it allows to analyze

graphs at different scales by looking at different diffusion times t . For small values of t , the diffusion only concerns a small neighborhood of the initially heated nodes, while for larger values it involves larger and possibly more complex structures, taking into account topological properties of the graph. Thus, the choice of relevant and informative diffusion times is essential.

For more references on comparisons of graphs, we refer the reader to [41], [12], [44] and references therein.

1.1 Our contributions

While a lot of the above notions of distances are often supported by experimental results and applications to learning or data mining tasks, they usually lack statistical foundations. In this context, we provide new tools to analyze and compare graphs or even data sets of graphs, that benefit from statistical guarantees. Our methods take advantage of the desirable multiscale property of heat diffusion. Moreover, one of our methods can deal with graphs without known NC or even graphs of different sizes, by using topological descriptors from topological data analysis (TDA). To circumvent the difficulty of choosing a suitable diffusion time, we opt to take into account the whole diffusion process. As a result, we define two real-valued processes, indexed by all the diffusion times in $[0, T]$ for some $T > 0$, representing comparisons of heat distributions. Basically, instead of trying to define a real value that would measure the distance between two graphs, we define a curve, a distance profile, that represents the comparisons of heat distributions across all diffusion times.

The first process, called Heat Kernel Distance (HKD) process, is defined by comparing heat kernels with the Frobenius norm. In this case, a NC between graphs needs to be known for the entry-wise comparison of the heat kernels to be meaningful. The second process, called Heat Persistence Distance (HPD), is defined using tools from TDA and can deal with graphs of different sizes. To do so, each graph is equipped with a real-valued function defined on the vertex set: the Heat Kernel Signature (HKS) [43, 23]. Then, graphs are converted into topological descriptors called persistence diagrams. They are multisets of points in \mathbb{R}^2 , encoding how topological features, like connected components and loops, evolve along with the families of sublevel and superlevel subgraphs. The diagrams are then compared with the so-called Bottleneck distance. Using persistence diagrams allows to switch from node-based representations of graphs to comparable topological summaries, hence requiring no assumption on graph sizes and NC.

To statistically study the HKD and HPD processes, we prove general results on Lipschitz-continuous real-valued empirical processes indexed by a real parameter. Namely, we show that they verify a functional Central Limit Theorem and admit Gaussian approximations with rates depending only on the sample sizes. These results ensure the asymptotic validity of bootstrap methods to design confidence bands around empirical mean processes, as well as consistent two-sample tests. They are applied to the HKD and HPD processes and could be applied to any other Lipschitz-continuous processes indexed by a real parameter under mild assumptions. In our graph framework, we can summarize the general idea as follows. We convert families of pairs of graphs into families of distance curves using HKD or HPD processes and then, we leverage the statistical properties of smooth curves mentioned above to exploit these families of curves. As a proof of concept, we illustrate these results on simulated data sets of pairs of graphs, drawn from various models: Erdős-Rényi model, stochastic block model, and random geometric graph models. We also compare our HKD-based test with tests based on other notions of distance and show improvements when using the multi-scale comparison of graphs. The HPD-based test is compared with various tests based on graph kernels.

The results from this article allow to statistically analyze data sets of pairs of graphs. While it might seem more natural to study data sets of graphs, we would like to motivate our analysis. There are various situations where graphs naturally come in pairs, and where the relevant information is actually the structural changes inside the pairs. For example, when monitoring brain diseases [37, 16], patients' brains can be observed at different time points. In this case, the changes of the brains' connectivity are of prime interest to assess the evolution of the disease [15, 6]. Moreover, even when graphs don't come in pairs, there exist several ways to turn a data set of graphs into a data set of pairs of graphs. For example, one can consider all the possible pairs of graphs from the data set or split the data set into two groups and construct pairs that contain a graph of each group. However, designing relevant approaches to combine graphs into pairs is probably a task-dependent problem, which is beyond the scope of this article.

1.2 Organisation

The rest of the paper is organized as follows. Section 2 introduces the graph framework, heat diffusion on graphs, the HKD and HPD processes. General results of such processes are developed in Section 3. We introduce the framework for general continuous real-valued empirical processes indexed by a real parameter,

prove their statistical properties and present some applications on bootstrap methods. These general results are applied to the HKD and HPD processes in Section 4, where we also detail the construction of a HKD or HPD-based confidence bands and two-sample tests for samples of pairs of graphs. We provide asymptotic results on the level and power of these tests. Finally, as a proof of concept, we illustrate the construction of confidence bands and two-sample tests in Section 5 using several generative models of random graphs. All Python codes are available at <https://github.com/elasalle/HeatDistanceProcess>.

2 Comparing graphs using heat distance processes

In this section, we introduce notions of distance between graphs based on heat diffusion. We define the resulting HKD and HPD processes.

2.1 Background and definitions

Before introducing the key notions of this work, we present general definitions and notations that will be used in the rest of the paper. We start by introducing notations relative to graph theory before presenting the theory of extended persistence adapted to graphs.

2.1.1 Graphs

For $0 \leq w_{\min} \leq w_{\max}$, we denote by $\mathcal{G}_n(w_{\min}, w_{\max})$ the set of undirected weighted graphs of size n , without self-loop and whose weights are in $\{0\} \cup [w_{\min}, w_{\max}]$. The special case of unweighted graphs correspond to $w_{\min} = w_{\max} = 1$. For clarity in the notation, we remove the w_{\min} and w_{\max} , whenever there is no ambiguity. We also consider \mathcal{G}^n (with n as an exponent) the set of graphs of size at most n , *i.e.*, $\mathcal{G}^n = \cup_{1 \leq i \leq n} \mathcal{G}_i$. For a graph G in \mathcal{G}_n , we denote by $W(G)$ its weight matrix (or adjacency matrix), *i.e.*, the $n \times n$ symmetric matrix whose (i, j) -coefficient is the weight $w_{i,j}$ of edge $\{i, j\}$. The degree matrix $D(G)$ denotes the diagonal matrix whose entry $D(G)_{i,i}$ is the degree of node i defined by $\sum_j w_{i,j}(G)$, the sum of all incident weights. The combinatorial Laplacian $L(G)$ is defined by $D(G) - W(G)$. Taking non-negative weights ensures that $L(G)$ is a real symmetric positive-semidefinite matrix. From now on, we forget the dependence in G in the notation, whenever there is no ambiguity. Let $\lambda_1 \leq \dots \leq \lambda_n$ be the eigenvalues of L and let (ϕ_1, \dots, ϕ_n) be a family of orthonormal eigenvectors. We denote by Λ the diagonal matrix containing the eigenvalues on the diagonal and ϕ the matrix whose columns are the ϕ_i 's so that L admits the following decomposition

$$L = \phi \Lambda \phi^T = \sum_{k=1}^n \lambda_k \phi_k \phi_k^T. \quad (2.1)$$

Note that $\lambda_1 = 0$ and that ϕ_1 can always be chosen to be the vector whose entries are equal to $1/\sqrt{n}$. In the following, this choice will always be made.

2.1.2 Persistence on graphs

We present here the basics of ordinary and extended persistence. We refer the reader to [8], [11] and [33] for a complete description of these theories.

Persistence theory allows to study the topology of topological spaces in a multiscale manner. Usually, given a topological space X and a continuous real-valued function $f : X \rightarrow \mathbb{R}$, one considers the family of sublevel sets $X_\alpha := \{x \in X, f(x) \leq \alpha\}$, for α varying from $-\infty$ to $+\infty$. Ordinary persistence records the levels at which topological features (connected components, loops, cavities or higher dimensional holes...) appear and disappear. For each feature, its birth and death levels are stored as the coordinates (b, d) of a point in \mathbb{R}^2 . The multiset of these points is called a persistence diagram. This framework can be applied to graphs.

Let $G = (V, E)$ be a graph with vertex set V and edge set E , and f be a real-valued function on V . Consider the family of sublevel subgraphs $(G_\alpha)_{\alpha \in \mathbb{R}}$, where $G_\alpha = (V_\alpha, E_\alpha)$ with $V_\alpha = \{v \in V, f(v) \leq \alpha\}$ and $E_\alpha = \{\{v, v'\} \in E, v, v' \in V_\alpha\}$. Across the family of sublevel graphs, as α increases, we can record birth and death levels of connected components, and birth levels of loops. As a connected component dies when it gets connected to an older connected component, remark that the connected components of G will never die. Similarly, as a loop dies when it gets "filled-in" by a 2-dimensional object, loops appearing in G_α (an object of maximal dimension 1) will never die. To prevent topological features from having no death levels (or infinite death levels), the theory of extended persistence suggests that the family of superlevel subgraphs

should also be considered. Define $G^\alpha = (V^\alpha, E^\alpha)$ similarly to G_α with $V^\alpha = \{v \in V, f(v) \geq \alpha\}$. A death level is now assigned to connected components of G and loops as the level at which they appear in the family of superlevel subgraphs when α decreases. Additionally, we record the birth and death of connected components in the family of superlevel subgraphs. Hence, extended persistence is able to detect four types of topological features and extract their birth and death levels (b, d) corresponding to four types of points:

- Ord_0 : birth and death of a connected component in (G_α) .
- Rel_1 : birth and death of a connected component in (G^α) .
- Ext_0 : birth and death of a connected component of G when using (G_α) and (G^α) .
- Ext_1 : birth and death of a loop when using both (G_α) and (G^α) .

These four types of topological features can be seen as downward branches, upward branches, connected components, and loops, respectively; with the orientation being taken with respect to f . For a graph G and a function f on its vertices, we will denote by $Ord_0(G, f)$, $Rel_1(G, f)$, $Ext_0(G, f)$ and $Ext_1(G, f)$ the persistence diagrams containing the corresponding points. In the following, $Dg(G, f)$ will generically denote any of these four diagrams. We refer the reader to [5, Section 2.1] for a more precise and illustrative presentation of extended persistence diagrams on graphs.

The space of diagrams can be equipped with the Bottleneck distance d_B . We recall its definition. Let μ and ν be two diagrams, *i.e.*, two multisets of points in \mathbb{R}^2 , and let $\Delta := \{(a, a), a \in \mathbb{R}\}$ be the diagonal. Denote by $\Pi(\mu, \nu)$ the set of bijections from $\mu \cup \Delta$ to $\nu \cup \Delta$. d_B is defined by

$$d_B(\mu, \nu) := \inf_{\pi \in \Pi(\mu, \nu)} \sup_{x \in \mu \cup \Delta} \|x - \pi(x)\|_\infty.$$

We state a stability result for extended persistence diagrams computed on graphs. It is a consequence of a more general stability result for persistence diagrams [7, 8].

Theorem 1. *For all graphs $G = (V, E)$, for all $f, f' : V \rightarrow \mathbb{R}$, and for all diagram constructions Dg among Ord_0, Rel_1, Ext_0 and Ext_1 ,*

$$d_B(Dg(G, f), Dg(G, f')) \leq \|f - f'\|_\infty,$$

with $\|f\|_\infty = \max\{|f(v)|, v \in V\}$.

2.2 Heat distances

Let G be a graph in \mathcal{G}_n and let L be its Laplacian. For $t \geq 0$, let $u_t \in \mathbb{R}^n$ be the vector whose i -th coefficient represents the amount of heat of node i at time t . Then, u_t follows the heat equation:

$$\forall t \geq 0, \quad \frac{d}{dt} u_t = -L u_t.$$

The solution is given by $u_t = e^{-tL} u_0$. The matrix e^{-tL} is called the *heat kernel* and describes how heat diffuses in the graph. The i -th column of e^{-tL} contains the amount of heat of each node at time t , when a single unit of heat was placed at node i at time $t = 0$. From (2.1), the heat kernel decomposes in

$$e^{-tL} = \phi e^{-t\Lambda} \phi^T = \sum_{k=1}^n e^{-t\lambda_k} \phi_k \phi_k^T.$$

The heat kernel encodes all the solutions of the heat equation and will be used to define notions of distances between graphs. Note that the diffusion time t acts as a scale parameter, where diffusion for small values of t only considers direct neighborhoods of the nodes, while for larger values it takes more global structures into account. This multi-scale property will be exploited in the following.

2.2.1 Heat Kernel Distance

Assume that we know the NC for graphs in \mathcal{G}_n and that we number the nodes such that the identity mapping gives the correspondences. Hence, comparing adjacency matrices, Laplacians, or heat kernels entry-wise becomes meaningful. Here we compare graphs through their heat kernels. For two graphs G and G' , define their Heat Kernel Distance (HKD) at time t by

$$D_t((G, G')) = \|e^{-tL} - e^{-tL'}\|_F, \tag{2.2}$$

Heat diffusion distance processes




Figure 1: Top: representation of three unweighted graphs G , G' , G'' . Removing the edge $\{5, 6\}$ to G gives G' . Removing $\{1, 2\}$ to G gives G'' . Bottom: we plot $t \mapsto D_t((G, G'))$ and $t \mapsto D_t((G, G''))$.

where L and L' are the laplacian matrices of G and G' respectively, and $\|\cdot\|_F$ denotes the Frobenius norm. This notion of distance was introduced by [20]. Basically, it compares different signals defined on the nodes of G and G' . These signals are given by the heat distribution obtained after a diffusion time t , for all elementary initial conditions. Considering heat kernels instead of adjacency or Laplacian matrices allows for a more robust notion of distance. Indeed, it is more aware of the global structures of graphs. Adding or removing edges with low impacts on the overall graphs structures results in small changes of the HKDs, especially for diffusion times that are not too close to zero. For example removing an edge joining two dense components is seen differently by the HKD than removing an edge inside one of the component (see Figure 1). This would not be the case when comparing adjacency or laplacian matrices. The example in Figure 1 also illustrates the multi-scale property of HKDs. At a local scale the HKD only detects that an edge has been removed. However, at a larger scale the functional role of the removed edge for the graph connectivity is captured by the HKD. Moreover, we choose to compare the heat kernels through the Frobenius norm as it is a Euclidean norm on the space of matrices. This property will be handy for deriving an exact expression of $D_t((G, G'))$ in terms of the eigen-elements of L and L' in Proposition 7.

To turn the HKD into a parameter-free notion of distance, [20] define the *Graph Diffusion Distance* as $\max_t D_t((G, G'))$. This has the drawback of comparing different pairs of graphs at different times. To avoid this effect and to take advantage of the multi-scale property of heat kernels mentioned above, our approach consists in using the whole function $t \rightarrow D_t((G, G'))$. More precisely, considering a probability distribution P on $\mathcal{G}_n \times \mathcal{G}_n$ and a random pair of graphs $(G, G') \sim P$, we are interested in the stochastic process $\{D_t((G, G')), t \in [0, T]\}$ for some $T > 0$. That is, the process obtained by evaluating the functions of the family $\mathcal{F}_{HKD} := \{D_t, t \in [0, T]\}$ on a random pair of graphs (G, G') . This framework corresponds to the general framework of *empirical processes*.

2.2.2 Heat Persistence Distance

In practice, the NC between graphs is not always known. Additionally one may be interested in comparing graphs of different sizes. In these cases, HKD cannot be computed. To circumvent these issues and following ideas from [5], we define the Heat Persistence Distance (HPD) by using extended persistence diagrams computed with the Heat Kernel Signature (HKS). These persistence diagrams can be compared with the Bottleneck distance d_B without any assumption on graph sizes and node correspondence.

The HKS was first introduced by [43] for the study of shapes. Here we restrict ourselves to the definition of the HKS on graphs of [23]. For a graph G of size n with vertex set $V = \{1, \dots, n\}$, the HKS at time t is the

function $h_t(G) : V \rightarrow \mathbb{R}$ such that

$$h_t(G)(i) = \sum_{k=1}^n e^{-t\lambda_k} \phi_k(i)^2, \quad 1 \leq i \leq s.$$

Intuitively, the image of $h_t(G)$ corresponds to the diagonal of the heat kernel e^{-tL} . Hence, $h_t(G)(i)$ represents the remaining amount of heat at node i after a diffusion time t , when a single unit of heat was placed at node i at time $t = 0$. For each value of t , the HKS provides a function on the vertices of a graph, that we use to compute extended persistence diagrams. Recall that these persistence diagrams encode the upward branches, downward branches, connected components and loops of the graph, when the up/down orientation is given by the HKS. The HPD at time t between two graphs G, G' in \mathcal{G}^n is defined by

$$H_t((G, G')) = \max_{Dg} d_B(Dg(G, h_t(G)), Dg(G', h_t(G'))), \quad (2.3)$$

where the maximum is taken over the four diagram constructions Ord_0, Rel_1, Ext_0 and Ext_1 . In our simulations, persistence diagrams are computed by following the approach of [5] and using the Gudhi library [30].

Similarly to \mathcal{F}_{HKD} , we define the family $\mathcal{F}_{HPD} := \{H_t, t \in [0, T]\}$ in order to study the induced stochastic process: $\{H_t((G, G')), t \in [0, T]\}$, for some random pair of graphs $(G, G') \in \mathcal{G}^n \times \mathcal{G}^n$.

The statistical properties of the processes associated with \mathcal{F}_{HKD} and \mathcal{F}_{HPD} are developed in Section 4.1. But we first provide in Section 3 a more general study of such empirical processes.

3 General continuous empirical processes

In this section, we properly introduce the general framework for continuous empirical processes, then show that uniform boundedness and Lipschitz-continuity implies a functional central limit theorem, as well as a Gaussian approximation. Finally, we derive consequences on the construction of confidence bands and two-sample tests.

3.1 Background and definitions

Let I be a compact interval of \mathbb{R} and $\mathcal{C}(I)$ the space of continuous real-valued functions on I endowed with the metric induced by the uniform norm: $\|h\|_\infty = \sup_{t \in I} |h(t)|$. Consider a measurable space $(\mathbb{X}, \mathcal{X})$. For all measures Q on $(\mathbb{X}, \mathcal{X})$ and all measurable functions $g : \mathbb{X} \rightarrow \mathbb{R}$, we denote the integral of g with respect to Q by $Qg := \int_{\mathbb{X}} g(x) dQ(x)$. Consider a probability measure P on $(\mathbb{X}, \mathcal{X})$ and $\mathcal{F} := \{f_t, t \in I\}$, a family of measurable real-valued functions on \mathbb{X} indexed by I . For all $x \in \mathbb{X}$, define $f(x)$ as the function $t \rightarrow f_t(x)$, and assume that $f(x) \in \mathcal{C}(I)$. Therefore, given a random variable X with distribution P , one can equivalently see $\{f_t(X), t \in I\}$ either as a random process or as $f(X)$ a random variable in $\mathcal{C}(I)$.

Given an i.i.d sample X_1, \dots, X_N drawn under P , we are interested in the statistical properties of the mean function $N^{-1} \sum_i f(X_i)$ and its centered and scaled version $N^{-1/2} (\sum_i f(X_i) - Pf)$. Equivalently, one can study the empirical processes $\{P_N f_t, t \in I\}$ and $\{G_N f_t, t \in I\}$, where $P_N = N^{-1} \sum_i \delta_{X_i}$, and $G_N = \sqrt{N}(P_N - P)$. In the following, we see random processes and random functions as the same objects.

When studying the statistical properties of \mathcal{F} , one might be interested in a functional version of the central limit theorem. This corresponds to the concept of Donsker families.

Definition 1 (P-Donsker). Let P be a probability measure on $(\mathbb{X}, \mathcal{X})$, the family \mathcal{F} is called P -Donsker if the process $\{G_N f_t, t \in I\}$ converges in distribution to the centered Gaussian process $\{\mathbb{G}_t, t \in I\}$ with covariance function κ defined as $\kappa_{s,t} := Pf_t f_s - Pf_t Pf_s$ for all $t, s \in I$.

Here convergence in distribution means weak convergence in the space $\mathcal{C}(I)$. That is, for all continuous bounded functions $h : \mathcal{C}(I) \rightarrow \mathbb{R}$, $\lim_{N \rightarrow \infty} \mathbb{E}[h(G_N f)] = \mathbb{E}[h(\mathbb{G})]$, where the expectation on the left-hand side is taken over the distribution of the sample X_1, \dots, X_N , and the one on the right-hand side is taken over the distribution of the Gaussian process \mathbb{G} .

Going further into the statistical analysis of \mathcal{F} , one might want to assess the speed at which $\{G_N f_t, t \in I\}$ converges in distribution to $\{\mathbb{G}_t, t \in I\}$. This can be done by proving Gaussian approximation results.

Definition 2 (Gaussian Approximation). Let $(r_N)_{N \geq 1}$ be a vanishing sequence of positive real numbers. We say that the process $\{G_N f_t, t \in I\}$ admits a *Gaussian approximation with rate r_N* , if for all $\lambda > 1$, there

exists a constant C such that for all $N \geq 1$ one can construct on the same probability space both the sample X_1, \dots, X_N and a version $\mathbb{G}^{(N)}$ of the Gaussian process \mathbb{G} verifying

$$\mathbb{P} \left(\left\| G_N f - \mathbb{G}^{(N)} \right\|_{\infty} > C r_N \right) \leq N^{-\lambda}.$$

Note that if $\{G_N f_t, t \in I\}$ admits a Gaussian approximation with rate r_N , applying the Borel-Cantelli Lemma would yield $\|G_N f - \mathbb{G}^{(N)}\|_{\infty} = O(r_N)$ *almost surely*.

3.2 Donsker theorem and Gaussian approximation

Before stating the central theorem of the paper, we introduce two assumptions on \mathcal{F} .

- (L) - There exists $k > 0$ such that for all $x \in \mathbb{X}$ the function $t \rightarrow f_t(x)$ is k -Lipschitz continuous on I , meaning that for all $t, s \in I$

$$|f_t(x) - f_s(x)| \leq k |t - s|.$$

- (B) - \mathcal{F} is uniformly bounded. That is, there exists a constant $M > 0$ such that for all $x \in \mathbb{X}$ and for all $t \in I$,

$$|f_t(x)| \leq M.$$

Remark that assumptions (L) and (B) are simple and can easily be checked for most processes. We now show that they are sufficient to obtain a Donsker theorem and a Gaussian approximation result.

Theorem 2. *Assume that \mathcal{F} verifies assumptions (L) and (B).*

Then for any probability measure P on $(\mathbb{X}, \mathcal{X})$, \mathcal{F} is P -Donsker and $\{G_N f_t, t \in I\}$ admits a Gaussian approximation with rate $r_N = N^{-1/7} \log N^{9/14}$.

This theorem provides a functional central limit theorem as well as some information about the rate of convergence. It allows us to derive in the next section more practical consequences. Namely, it validates the construction of consistent confidence bands and consistent two-sample tests using bootstrap methods.

The proof of Theorem 2 can be found in Appendix A. The Donsker property is proved by standard arguments of tightness. The proof of Gaussian approximation is based on a result from [2] requiring technical work and a more complex formalism. Essentially, by using Lipschitz-continuity we control the covering number of \mathcal{F} , which is a quantity that indicates the complexity of the family of functions. Even if Theorem 2 is a consequence of [2], working in the special case of continuous processes indexed by a real parameter allows us to present a more accessible result. And although the main interest of the paper is its application to the HKD and HPD processes, we believe that the framework is simpler and that this theorem could easily be applied to other processes.

3.3 Statistical applications

Here we present how we can construct confidence bands around empirical mean processes and two-sample tests while retrieving statistical guarantees from Theorem 2.

Confidence Band

Let $c_{\alpha} := \inf\{u, \mathbb{P}(\|\mathbb{G}\|_{\infty} > u) \leq \alpha\}$ be the upper α -quantile of the maximum of the Gaussian limit process. As a consequence of Theorem 2, we have

$$\lim_{N \rightarrow \infty} \mathbb{P} \left(\forall t, P f_t \in \left[P_N f_t - \frac{c_{\alpha}}{\sqrt{N}}, P_N f_t + \frac{c_{\alpha}}{\sqrt{N}} \right] \right) \geq 1 - \alpha.$$

Unfortunately, as the distribution of \mathbb{G} is unknown, c_{α} can not be directly computed. Instead, consider a bootstrap sample $\hat{X}_1, \dots, \hat{X}_N$ drawn under P_N and let \hat{P}_N be its empirical probability measure. Consider the process $\{\hat{G}_N f_t, t \in I\}$ where \hat{G}_N is the measure $\sqrt{N}(\hat{P}_N - P_N)$ and let \hat{c}_{α} be the upper α -quantile of $\|\hat{G}_N f\|_{\infty}$ given the data. Theorem 2.6 in [25] ensures that when \mathcal{F} is P -Donsker, $\{\hat{G}_N f_t, t \in I\}$ converges weakly to \mathbb{G} , given the data. Hence \hat{c}_{α} provide an approximation of c_{α} , and we can use \hat{c}_{α} to design a consistent confidence band around $P_N f$:

$$\lim_{N \rightarrow \infty} \mathbb{P} \left(\forall t, P f_t \in \left[P_N f_t - \frac{\hat{c}_{\alpha}}{\sqrt{N}}, P_N f_t + \frac{\hat{c}_{\alpha}}{\sqrt{N}} \right] \right) \geq 1 - \alpha.$$

Note that \hat{c}_α can be estimated with Monte-Carlo simulations, by drawing as many bootstrap samples as we want.

Two Sample Test

Consider the following setup. Let P and Q be two probability distributions on \mathbb{X} , and assume we are given two independent iid samples (X_1, \dots, X_M) and (Y_1, \dots, Y_N) , drawn under P and Q , respectively. We denote by P_M and Q_N the empirical measures. We would like to test the null hypothesis $H_0 : P = Q$ against the alternatives $H_1 : P \neq Q$, by using the family \mathcal{F} , assuming that it is Donsker with respect to both P and Q . We follow the approach described in Section 3.7 of [46], and consider the following test statistic:

$$D_{M,N} := \sqrt{\frac{MN}{M+N}} \|P_M f - Q_N f\|_\infty.$$

The strategy is to define a data-dependent threshold $\hat{c}_{M,N}(\alpha)$ and reject the null hypothesis whenever $D_{M,N} > \hat{c}_{M,N}(\alpha)$. Consider the pooled data $(Z_1, \dots, Z_{M+N}) = (X_1, \dots, X_M, Y_1, \dots, Y_N)$, and its empirical measure H_{M+N} . Let $(\hat{Z}_1, \dots, \hat{Z}_{M+N})$ be a bootstrap sample drawn from H_{M+N} and consider the bootstrap empirical measures

$$\hat{P}_M = \frac{1}{M} \sum_{i=1}^M \delta_{\hat{Z}_i} \quad \text{and} \quad \hat{Q}_N = \frac{1}{N} \sum_{i=1}^N \delta_{\hat{Z}_{M+i}}.$$

We can define

$$\hat{D}_{M,N} := \sqrt{\frac{MN}{M+N}} \|\hat{P}_M f - \hat{Q}_N f\|_\infty,$$

as well as

$$\hat{c}_{M,N}(\alpha) = \inf \left\{ t, \mathbb{P} \left(\hat{D}_{M,N} > t \mid Z_1, \dots, Z_{M+N} \right) \leq \alpha \right\},$$

for $\alpha \in (0, 1)$. Note that $\hat{c}_{M,N}(\alpha)$ can be estimated with Monte-Carlo simulations. Using $\hat{c}_{M,N}(\alpha)$ as the threshold to accept or reject H_0 leads to a consistent test.

Theorem 3 ([46], Section 3.7.2). *Assume that \mathcal{F} is Donsker with respect to both P and Q and that $\|Pf\|_\infty$ and $\|Qf\|_\infty$ are finite. Furthermore assume that $M/(M+N) \rightarrow \lambda \in (0, 1)$. Then the test that rejects H_0 whenever $D_{M,N} > \hat{c}_{M,N}(\alpha)$ is consistent, in the sense that the asymptotic level is α and under any alternative verifying $\|Pf - Qf\|_\infty > 0$, $\mathbb{P}(D_{M,N} > \hat{c}_{M,N}(\alpha)) \rightarrow 1$.*

This provides guarantees of the two-sample test when both sample sizes grow to infinity. Intuitively, under H_0 , as both the process defining $D_{M,N}$ and the one defining $\hat{D}_{M,N}$ given Z_1, \dots, Z_{M+N} converge to the same Gaussian process, we can sample from the distribution of the latter (randomness coming only from resampling) to estimate the quantile of the former.

4 Back to the heat distance processes

Using the results of the previous section, we are able to show that the HKD and HPD processes admit a functional version of the central limit theorem, as well as a Gaussian approximation. This is essentially based on the fact that they are uniformly bounded Lipschitz-continuous processes. From these results, we propose a procedure to construct consistent confidence bands. We also detail the construction of a two-sample test for samples of pairs of graphs based on either the HPD processes or the HKD processes when the graph sizes and NC allow their computation. We explicit these constructions and state their consistency results.

4.1 Statistical properties

Let P be a probability distribution on $\mathcal{G}_n \times \mathcal{G}_n$. Let T be a positive real number and recall that $\mathcal{F}_{HKD} := \{D_t, t \in [0, T]\}$, with D_t defined in (2.2). We will study the centered and rescaled empirical process $\{G_N D_t, t \in I\} = \{\sqrt{N}(P_N - P)D_t, t \in I\}$, where P_N is the empirical measure $N^{-1} \sum_i \delta_{(G_i, G'_i)}$ associated to a N -sample $((G_1, G'_1), \dots, (G_N, G'_N))$ drawn under P . We now state the statistical results concerning the HKD process.

Theorem 4. For any probability distribution P on $\mathcal{G}_n \times \mathcal{G}_n$, the family \mathcal{F}_{HKD} is P -Donsker. That is, the process $\{\sqrt{N}(P_N - P)D_t, t \in [0, T]\}$ converges weakly to \mathbb{G} in $\mathcal{C}([0, T])$, where $\mathbb{G} = \{\mathbb{G}_t, t \in [0, T]\}$ is a zero mean Gaussian process with covariance function $\kappa(t, s) = P(D_t D_s) - P D_t \cdot P D_s$. Moreover, it admits a Gaussian approximation with rate $r_N = N^{-1/7} \log N^{9/14}$.

We have a similar result for HPD processes. Recall that for HPD processes we work in \mathcal{G}^n . Let P be a probability distribution on $\mathcal{G}^n \times \mathcal{G}^n$. Let T be a positive real number and recall that $\mathcal{F}_{HPD} := \{H_t, t \in [0, T]\}$, with H_t defined in (2.3). Again, we study the empirical process $\{G_N H_t, t \in I\} = \{\sqrt{N}(P_N - P)H_t, t \in I\}$, where P_N is the empirical measure $N^{-1} \sum_i \delta_{(G_i, G'_i)}$ associated to a N -sample $((G_1, G'_1), \dots, (G_N, G'_N))$ drawn under P .

Theorem 5. For any probability distribution P on $\mathcal{G}^n \times \mathcal{G}^n$, the family \mathcal{F}_{HPD} is P -Donsker. Thus, the process $\{\sqrt{N}(P_N - P)H_t, t \in [0, T]\}$ converges weakly to \mathbb{G} in $\mathcal{C}([0, T])$, where $\mathbb{G} = \{\mathbb{G}_t, t \in [0, T]\}$ is a zero mean Gaussian process with covariance function $\kappa(t, s) = P(H_t H_s) - P H_t \cdot P H_s$. Moreover, it admits a Gaussian approximation with rate $r_N = N^{-1/7} \log N^{9/14}$.

These theorems can be seen as functional central limit theorems for the HKD and HPD processes. They validate bootstrap methods for the construction of consistent confidence bands and consistent two-sample tests. Gaussian approximations strengthen these results and provide information about the speed of convergence to Gaussian processes. Note that this rate is independent of the graph size n . This could indicate that even in the finite sample case with large graphs, methods based on the Donsker property could work well.

The proof of Theorem 4 and Theorem 5 can be found in Appendix B. We start by showing that under the hypothesis of bounded weights on the edges of the graphs the HKD and HPD process are uniformly bounded and Lipschitz-continuous, then the proofs are direct applications of Theorem 2.

4.2 Confidence band and consistency

As explained in Section 3.3, Donsker results allow for the construction of consistent confidence bands using bootstrap methods. In Algorithm 1, we detail this construction based on the HKD and HPD processes on a sample of pairs of graphs. As the construction is similar in the HKD or HPD case, we only present the HKD case.

Algorithm 1: Confidence Band

Data: a sample of pairs of graphs $((G_1, G'_1), \dots, (G_N, G'_N))$

Input: an integer B and a level α .

for $i \in \{1, \dots, N\}$ **do**

 | compute $d_i : t \mapsto D_t((G_i, G'_i))$;

end

compute $\bar{d}_N = N^{-1} \sum_i d_i$;

for $b \in \{1, \dots, B\}$ **do**

 | draw with replacement a bootstrap sample $(\hat{d}_1^{(b)}, \dots, \hat{d}_N^{(b)})$ out of (d_1, \dots, d_N) ;

 | compute $\tilde{d}_N^{(b)} = N^{-1} \sum_i \hat{d}_i^{(b)}$;

 | compute $t^{(b)} = \sqrt{N} \|\tilde{d}_N^{(b)} - \bar{d}_N\|_\infty$;

end

compute \tilde{c}_α the empirical upper α -quantile of $t^{(1)}, \dots, t^{(B)}$;

Result: the confidence band $\left[\bar{d}_N(t) - \frac{\tilde{c}_\alpha}{\sqrt{N}}, \bar{d}_N(t) + \frac{\tilde{c}_\alpha}{\sqrt{N}} \right]$ for all t .

By choosing the bootstrap sample size B large enough, \tilde{c}_α can approximate as closely as we want the upper α -quantile \hat{c}_α of the bootstrap test statistic $t^{(1)}$ given the data. And according to Section 3.3 we know that

$$\lim_{N \rightarrow \infty} \mathbb{P} \left(\forall t, \mathbb{E}_P [D_t((G, G'))] \in \left[\bar{d}_N(t) - \frac{\hat{c}_\alpha}{\sqrt{N}}, \bar{d}_N(t) + \frac{\hat{c}_\alpha}{\sqrt{N}} \right] \right) \geq 1 - \alpha.$$

This means that for large enough B , Algorithm 1 produces a consistent confidence band.

4.3 Two-sample test and consistency

We now detail how to construct a two-sample test for sample of pairs of graphs in Algorithm 2. As previously, we present the HKD case, as the HPD case is similar.

Algorithm 2: Two-sample Test

Data: two samples of pairs of graphs

$((G_{1,1}, G'_{1,1}), \dots, (G_{1,N}, G'_{1,N}))$ and $((G_{2,1}, G'_{2,1}), \dots, (G_{2,M}, G'_{2,M}))$

Input: an integer B and a level α .

for $i \in \{1, \dots, N\}$ **do**

 | compute $d_{1,i} : t \mapsto D_t((G_{1,i}, G'_{1,i}))$;

end

for $i \in \{1, \dots, M\}$ **do**

 | compute $d_{2,i} : t \mapsto D_t((G_{2,i}, G'_{2,i}))$;

end

compute the empirical means $\bar{d}_1 = N^{-1} \sum_i d_{1,i}$ and $\bar{d}_2 = M^{-1} \sum_i d_{2,i}$;

compute the test statistic $T = \sqrt{NM(N+M)^{-1}} \|d_1 - d_2\|_\infty$;

for $b \in \{1, \dots, B\}$ **do**

 | draw with replacement $(\hat{d}_{1,1}^{(b)}, \dots, \hat{d}_{1,N}^{(b)})$ and $(\hat{d}_{2,1}^{(b)}, \dots, \hat{d}_{2,M}^{(b)})$ out of $(d_{1,1}, \dots, d_{1,N}, d_{2,1}, \dots, d_{2,M})$;

 | compute $\tilde{d}_1^{(b)} = N^{-1} \sum_i \hat{d}_{1,i}^{(b)}$ and $\tilde{d}_2^{(b)} = M^{-1} \sum_i \hat{d}_{2,i}^{(b)}$;

 | compute $T^{(b)} = \sqrt{NM(N+M)^{-1}} \|\tilde{d}_1^{(b)} - \tilde{d}_2^{(b)}\|_\infty$;

end

compute \tilde{c}_α the empirical upper α -quantile of $T^{(1)}, \dots, T^{(B)}$;

Result: Reject the null hypothesis if $T > \tilde{c}_\alpha$.

Once again, for a large enough B , \tilde{c}_α closely approximate the the upper α -quantile of $T^{(1)}$ given the data. Recall that from Theorem 3, using \hat{c}_α as the rejection threshold yields to a control of the asymptotic level and power. That is

$$\lim \mathbb{P}_{H_0} (T > \hat{c}_\alpha) \leq \alpha,$$

and under any alternative verifying $\|P_1 D. - P_2 D.\|_\infty > 0$ where P_1 and P_2 are the distributions that generated the two samples,

$$\lim \mathbb{P}_{H_1} (T > \hat{c}_\alpha) = 1.$$

The limits are for both N and M going to infinity under the condition that $N/(N+M) \rightarrow \lambda$, for a λ in $(0, 1)$. Hence, for large enough B , Algorithm 2 produces a asymptotically valid two-sample test for samples of pairs of graphs.

5 Experiments

We illustrate the construction of confidence bands and two-sample tests on synthetic data sets of pairs of graphs. For that, we consider different random graph models and combine them to create independent pairs of graphs.

5.1 Random graph models

In this section, we present the models generating the random graphs, namely the Erdős-Rényi model [13], the stochastic block model [22], the geometric model [34] and the Watts-Strogatz model [47]. For each, we specify the parameters used in our simulations.

Erdős-Rényi model (ER)

This model generates random graphs where each edge appears with probability p , independently from all the others. It requires two parameters: n the graph size and p the edge probability. Because of the independence and their homogeneity, ER graphs are considered to have no structure.

In our simulations, we take $n = 50$ and $p = 0.5$. Weights may be added by assigning a uniform weight between 0 and 2 to each existing edge, independently from all the others.

Stochastic block model (SBM)

This model is a generalization of the ER model that introduces a block structure. The n nodes are clustered in K groups C_1, \dots, C_K , of respective sizes n_1, \dots, n_K . Edges appear independently from the others, but with a probability depending on the groups: edge $\{i, j\}$ appears with probability $p_{k,l}$ when $i \in C_k$ and $j \in C_l$.

In our simulations, we take $K = 2$, $n_1 = n_2 = 25$, and $p_{1,1} = p_{2,2} = 0.75$, $p_{1,2} = p_{2,1} = 0.25$. So graphs are composed of two dense clusters, with few edges between them. Similarly to the ER model, we may add random weights following the uniform distribution between 0 and 2.

Geometric model (GM)

Given a compact domain U of \mathbb{R}^d for some d , a graph is generated from the GM by drawing n points uniformly on U and creating an edge between two points if their Euclidean distance is smaller than a given threshold. Here we choose a slight variation of this model by considering a number $p \in [0, 1]$ and creating the edges corresponding to the $\lfloor p \binom{n}{2} \rfloor$ pairs of points with the smallest euclidean distances.

In our simulations, the compact domain U is either A_ε the annulus in \mathbb{R}^2 with outer radius 1 and inner radius $\varepsilon > 0$ or A_0 the unit disk. We either fix $n = 50$ or for each graph, n is drawn from a Poisson distribution of parameter 50. We take $p = 0.5$. To obtain weighted graphs, we may assign the weight e^{-2d} to an edge, where d is the distance between the two points forming the edge.

Watts-Strogatz model (WS)

This model is famous for its so-called *small-world* property, meaning that any two nodes in the graph are usually not far apart in graph distance. Random graphs under this model are constructed as follows. First a ring of n nodes is created where each node is connected to its k nearest neighbors. Then, for each edge, with probability p independently of all others edges, one of its terminal nodes is replaced by another node chosen uniformly at random.

In our simulations, we take $n = 50$, $k = 25$ and $p = 0.5$. Weights may be added by assigning a uniform weight between 0 and 2 to each existing edge, independently from all the others.

We combine these models to generate pairs of independent graphs on which we can compute HKD and/or HPD processes. We consider the pairs of independent ER graphs (ER-ER) and the pairs containing one ER graph and one SBM graph (ER-SBM). For these distributions, the groups' composition is known and nodes are treated independently among groups. Thus, we can consider that we know an NC between graphs. As a result, we can compute both HKD and HPD processes. Similarly, we consider (ER-ER) and (ER-SW) pairs and compare them with HKD processes (as nodes are exchangeable). We also consider pairs of independent geometric random graphs: Disk-Disk and Disk-Annulus. However, as nodes in these models correspond to random points, there is no default NC. Hence, only HPD processes are computed.

5.2 Simulation results

In this section, we present the simulation results concerning the construction of confidence bands and the performances of the two-sample tests.

5.2.1 Confidence bands

In this section, we compute confidence bands under the different models of pairs of graphs defined above. For each model, we draw a sample $(G_1, G'_1), \dots, (G_N, G'_N)$ with $N = 100$. We compute the mean process, that is $t \rightarrow N^{-1} \sum_i D_t((G_i, G'_i))$ or $t \rightarrow N^{-1} \sum_i H_t((G_i, G'_i))$ and compute a confidence band of level 99% around this empirical mean using the bootstrap method detailed in Algorithm 1. Computations are done with 1000 bootstrap samples. Results are shown in Figure 2, 3 and 4, where solid lines represent empirical means and transparent areas represent confidence bands.

Remark that confidence bands around HKD processes (Figure 2) seem to be narrower than those around HPD processes (Figure 3), relatively to the height of the curves. Therefore, users should rather use HKD processes whenever NC's are available. Nonetheless, the versatility of HPD processes does not totally reduce

Heat diffusion distance processes




Figure 2: Confidence band around the mean HKD processes with ER-ER (red) and ER-SBM (blue) distributions.




Figure 3: Confidence band around the mean HPD processes with unweighted ER-ER (red) and ER-SBM (blue) distributions.

their efficiency. As Figure 4 indicates, HPD empirical means seem to be able to discriminate between the different distributions. These observations will be confirmed in the next section, where the performances of the two-sample tests are investigated.

5.2.2 Two-sample tests

Levels and Powers. Simulations are run to evaluate the performances of the two-sample tests using HKD and HPD processes as detailed in Algorithm 2, see Figure 5, Figure 6 and Figure 7. In all tests, the desired level is set to 0.05, and computations are done with 1000 bootstrap samples. Figures 5a and 6a illustrate that the asymptotic levels of the tests correspond to the set level. On the other hand, Figures 5b, 6b and 7 illustrate




Figure 4: Confidence band around the mean HPD processes with Disk-Disk (red) and Disk-Annulus (blue) distributions.

Heat diffusion distance processes




Figure 5: Performances of the two-sample test using HKD processes. The level estimations are made with weighted ER-ER distributions, while the power estimations are made with weighted ER-ER and ER-SBM distributions. The samples size varies in [25,50,100,200]. Estimations are done by repeating 400 independent tests. Vertical lines represent 95%-confidence intervals of the estimations.




Figure 6: Performances of the two-sample test using HPD processes. The level estimations are made with Disk-Disk distributions, while the power estimations are made with Disk-Disk and Disk-Annulus distributions. Graphs are unweighted and with fixed size. The samples size varies in [20,40,80,120]. Estimations are done by repeating 100 independent tests. Vertical lines represent 95%-confidence intervals of the estimations.

that the powers tend to 1 when sample sizes increase, indicating that the tests manage to distinguish between the different distributions. Moreover, Figure 6b and 7 show the good performances of the HPD-based test even when NC is unknown and when graphs have different sizes.

Comparison with other distances. We first propose to compare our HKD-based two-sample test with tests based on other distances. We consider the *Graph Diffusion Distance* (GDD) of [20] and the comparison of graphs based on the Frobenius norm of the difference of their adjacency matrices. As these distances return real values and not processes anymore, we compare samples of pairs of graphs through the mean distance of each sample and compute a rejection threshold using a standard bootstrap method. The result of these comparisons are presented in Figure 8. In Figure 8a, we consider distributions of pairs of ER-ER vs pairs of ER-SBM. In that case, both the adjacency-based and GDD-based test are unable to detect that graph distributions are different, while the HKD-based test reaches a power of 1 for samples of size greater than 100. In Figure 8b, we evaluate the tests on pairs of ER-ER vs pairs of ER-WS. In this setting, only the adjacency-based test fails to distinguish between the two distribution, while the HKD-based and GDD-based test have similar powers.

It is particularly interesting to witness the power difference (and similarity) between the tests with GDD and HKD. Essentially, the GDD quickly discards the multi-scale information by taking the maximum of each HKD process. On the other hand, the HKD-based test conserves it until it compares the two mean processes. Figure 8a illustrates that the multi-scale information contained in the HKD-processes can actually be relevant, and that it is worth conserving it to compare samples of graphs. However, from Figure 8b,

Heat diffusion distance processes




Figure 7: Power of the two-sample test using HPD processes with graphs of random sizes. Power estimations are made with Disk-Disk and Disk-Annulus distributions. Graphs are unweighted and with Poissonian sizes. The samples size varies in [25,50,100,150]. Estimations are done by repeating 100 independent tests. Vertical lines represent 95%-confidence intervals of the estimations.

we see that this additional multi-scale information is not always needed. A possible explanation could be that when considering SBM graphs, the density of edges is heterogeneous. For example, the density of edges in the whole graph is different compared to the one inside each block. Thus, heat diffusion may reflect that heterogeneity. To give the intuition, at the beginning of the diffusion process, heat essentially diffuses inside a dense block, then it diffuses to the rest of the graph more slowly as the edge density is smaller. Thus, capturing these two phases of the heat diffusion can be meaningful. However, in the ER and WS models, graphs are essentially homogeneous. In this case, multi-scale information might be less valuable.

We then proceed by comparing our HPD-based two-sample test with tests based on graph kernels and the *Maximum Mean Discrepancy* (MMD) [19]. We apply them on graphs with different sizes. We consider the

- GS - graphlet sampling [39]
 - WL - Weisfeiler-Lehman [38]
 - SP - shortest path [4]
 - VH - vertex histogram [42]
 - ST - SVM Theta [24]
 - NH - neighborhood Hash [21]
 - RW - random walk [42].
- following graph kernels: Kernel computations were done using

the Grakel python library [40].

In this set of experiments, we only consider unweighted graphs as most of these kernels do not handle edge weights. For kernels requiring vertex labels, we choose to provide the degree of a vertex as its label. Moreover, MMD-based tests essentially perform a test on the mean of each sample seen in the corresponding RKHS. Therefore, they are not designed to handle sample of pairs of graphs. As a consequence let us give a description of the experimental procedure. We want to evaluate how well two-sample tests distinguish between two different graph distributions P and Q . HPD-based test will be applied to samples of independent pairs P - P and P - Q while MMD-based tests are computed with samples of graphs from P and from Q . Figure 9a displays the powers the various tests for P being the Disk distribution and Q the Annulus distribution, each with Poissonian graph sizes. The HPD-based test achieve good performance for sample sizes greater than 100 pairs of graphs, but it is out-performed by the tests based on the graphlet sampling, Weisfeiler-Lehman, vertex histogram and SVM Theta kernels. Moreover, the three other kernel-based tests are unable to distinguish between P and Q . Note that among the successful kernels, the WF and ST require labels and thus their tests are actually based on degree distributions, which are different in P and Q . On Figure 9b, the test powers are given for P being the ER distribution and Q the SBM one, each with Poissonian graph sizes. Here, only the GS and ST methods successfully distinguish between P and Q , all other methods fail.

Comparison with the Neyman-Pearson Test. The Neyman-Pearson test [32] is an optimal testing procedure, in the sense that it is the test with the highest power for a given level. But being based on likelihood ratios, it rarely is computable. Its performances are determined by the total variation distance (TV) between the two distributions. If we consider the case of distributions that depend on the sample size, the speed at which

Heat diffusion distance processes

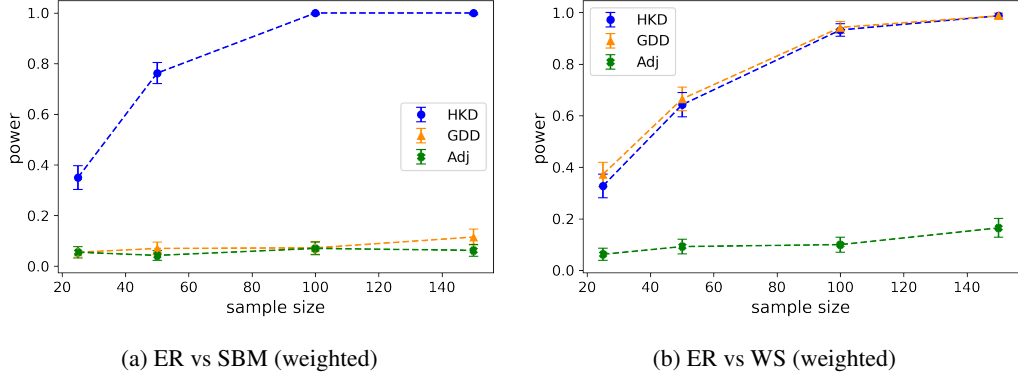


Figure 8: Powers of the HKD-based test, the GDD-based test and the adjacency-based test. Tests are performed either under weighted ER-ER and ER-SBM distributions (left) or weighted ER-ER and ER-WS distributions (right). The samples size varies in $[25,50,100,150]$. Estimations are done by repeating 400 independent tests. Vertical lines represent 95%-confidence intervals of the estimations.

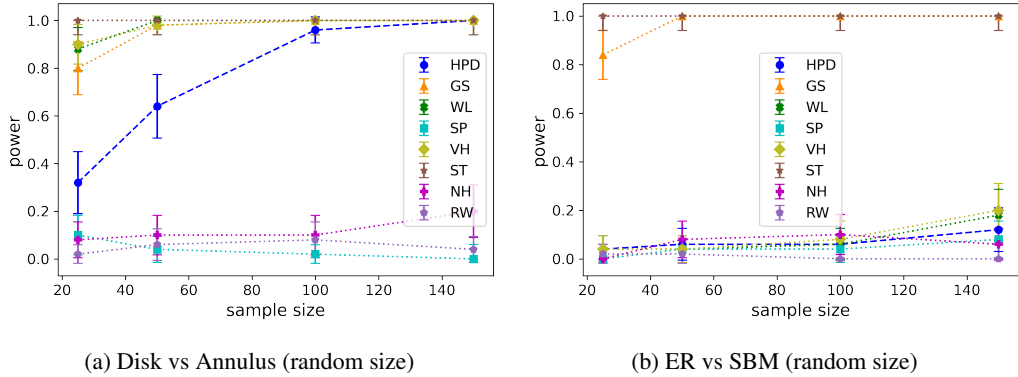


Figure 9: Powers of the HPD-based test and MMD tests for various graph kernels. Tests are performed either under Disk-Disk and Disk-Annulus distributions with graphs of random sizes (left) or ER-ER and ER-SBM distributions with graphs of random sizes (right). The samples size varies in $[25,50,100,150]$. Estimations are done by repeating 50 independent tests. Vertical lines represent 95%-confidence intervals of the estimations.

their TV tends to 0 determines if the Neyman-Pearson test will asymptotically be able to distinguish between the two distributions. For ER models, independence of the edges allows to theoretically compute bounds of the TV and determine the phase transition. Let us take a real number p such that $0 < p < 1$. And for each sample size $N \geq 1$ consider the parameters $p_0(N)$ and $p_1(N)$, such that both converge to p . Following [28, Section 13.1.], we can show that as long as $|p_0(N) - p_1(N)| \gg N^{-1/2}$, the Neyman-Pearson test asymptotically distinguishes between the distributions of independent pairs of $ER(n, p_0(N))$ and independent pairs of $ER(n, p_1(N))$ when using N -samples. Figure 10 shows that for large enough sample sizes, our two-sample test based on HKD processes distinguishes between ER models up to $|p_0(N) - p_1(N)| = C \log N / \sqrt{N}$, with $C = 0.01$. The tests are computed with a set level of 0.05 and with 1000 bootstrap samples.

5.3 Computational complexity

Both Algorithm 1 and Algorithm 2 require to compute eigenvalues and eigenvectors decompositions of Laplacian matrices in order to obtain heat diffusion distance processes. For symmetric matrices of size $n \times n$, the computational complexity of *divide-and-conquer* algorithms to compute such decomposition is, in the worst case, in $O(n^3)$ operations. (see Table 3.1 of [31]).

In our experiments, we restricted graph sizes to around 50 vertices. This is because we needed to repeat the tests several hundred times to estimate the levels and powers with statistical significance. However, single tests with graphs of a few hundreds or even thousands of vertices could be considered.




Figure 10: Evolution of the power of the HKD two-sample test when $|p_0(N) - p_1(N)| = C \log N / \sqrt{N}$. The samples size N varies in $[20, 50, 100, 200, 300]$. Estimations are done by repeating 100 independent tests. Vertical lines represent 95%-confidence power intervals of the estimations.

In the case where graphs would be too large for eigendecompositions, strategies to approximate the heat kernel could be developed. Previous works used Taylor expansion, Chebychev approximation, or low frequency approximation [45, 29].

6 Conclusion

We proposed two multiscale comparisons of graphs using heat diffusion processes, namely the HKD and HPD. The first one requires the assumption of equal graph sizes and a known NC, while the second one is free of these assumptions. The multiscale approach solves the problem of choosing an informative diffusion time. We proposed to use these processes to analyze data sets of pairs of graphs and were able to design consistent confidence bands and two-sample tests. The methods are supported by theoretical results: the HKD and HPD families are Donsker, meaning that the processes verify a functional central limit theorem. Moreover, the processes admit Gaussian approximations with rates that are independent of the graph sizes. These results are very general and could easily be applied to other processes. Essentially, the processes are required to be uniformly bounded and Lipschitz-continuous. Moreover, the performances of our methods were illustrated by simulations on synthetic data sets. We showed that the two-sample tests were able to distinguish between Erdős-Rényi and SBM graphs, as well as between geometric graphs sampled on different domains. Moreover, we confirmed the benefit of using multi-scale notion of distances, as the HKD and HPD processes, over other notions of distances, especially when graphs have heterogeneous densities. On Erdős-Rényi models with parameters depending on the sample size, the tests were still distinguishing between the different distributions, even when working close to the phase transition of the Neyman-Pearson test.

As future work, we would like to extend these methods to be able to perform learning tasks, e.g., clustering, classification, or change point detection on graph data. Extensions to be able to deal with data sets of graphs by opposition to pairs of graphs should also be developed to broaden the application spectrum. Transferring our approach to directed graphs could also be interesting. Although technical difficulties related to spectral decomposition of non-symmetric matrices would need to be dealt with. On the theoretical side, studying the interplay between graph sizes and sample sizes could be a first step toward non-asymptotic methods to analyze data sets of graphs.

Nonetheless, we believe that the introduction of the HKD and HPD processes has the potential to bring innovative and statistically founded ways to analyze data sets of graphs. Moreover, the theoretical results presented in this work being very general, our methods could be extended to other fields.

Appendix

A Proof of Theorem 2

We start by proving the Donsker property.

Proof. We need to prove the weak convergence of $\{G_N f_t, t \in I\}$ to the centered gaussian process \mathbb{G} . To do so, remark that assumption **(B)** ensures that second moments of $f(X)$ are finite. This allows to apply the multidimensional version of the central limit theorem to all finite-dimensional marginals of the random function $G_N f$. Hence, these finite-dimensional marginals converge in distribution to those of \mathbb{G} . To conclude to the weak convergence of the whole process, the sequence of random functions $\{G_N f, N \geq 1\}$ needs to be tight. According to [3, Theorem 12.3], tightness of $\{G_N f, N \geq 1\}$ is implied by the following two conditions:

1. There exists $t_0 \in I$, such that $\{G_N f_{t_0}, N \geq 1\}$ is tight.
2. There exist $\gamma \geq 0, \alpha > 1$ and a non-decreasing function $\psi : I \rightarrow \mathbb{R}$ such that $\forall t, s \in I, \forall N \geq 1$,

$$\mathbb{E} [|G_N f_t - G_N f_s|^\gamma] \leq |\psi(t) - \psi(s)|^\alpha.$$

Let us start by proving point 1. Let t_0 be any point in I . We have to prove that for all $\eta > 0$, there exists $\alpha > 0$ such that for all $N \geq 1$

$$P(|G_N f_{t_0}| \leq \alpha) > 1 - \eta. \quad (\text{A.1})$$

If $\text{Var}(f_{t_0}(X)) = 0$, $\{G_N f_{t_0}, N \geq 1\}$ is tight, since for all N , $G_N f_{t_0} = 0$ P -a.s.. Otherwise, fix $\eta > 0$. As the left hand side in (A.1) is non-decreasing with respect to α , we may as well show that there exists α such that for all N , $P(-\alpha < G_N f_{t_0} \leq \alpha) > 1 - \eta$. Following from **(B)**, $f_{t_0}(X)$ admits a third moment. Combined with the positive variance, we can apply the Berry-Essen Theorem: if F_N and ϕ denote the cumulative distribution functions of, respectively, $G_N f_{t_0}$ and a centered Gaussian variable of variance $\text{Var}(f_{t_0}(X))$, then there exists a constant C such that for all $\alpha \in \mathbb{R}$ and all $N \geq 1$

$$|F_N(\alpha) - \phi(\alpha)| \leq \frac{C}{\sqrt{N}}.$$

Now we take N_η such that for all $N > N_\eta$, $C/\sqrt{N} < \eta/4$, and we choose α_η such that $\phi(\alpha_\eta) > 1 - \eta/4$. Then, for all $N > N_\eta$

$$\begin{aligned} & P(-\alpha_\eta < G_N f_{t_0} \leq \alpha_\eta) \\ &= F_N(\alpha_\eta) - F_N(-\alpha_\eta) \\ &= \phi(\alpha_\eta) - \phi(-\alpha_\eta) + (F_N(\alpha_\eta) - \phi(\alpha_\eta)) - (F_N(-\alpha_\eta) - \phi(-\alpha_\eta)) \\ &\geq \phi(\alpha_\eta) - \phi(-\alpha_\eta) - 2\frac{C}{\sqrt{N}} \\ &= 2\phi(\alpha_\eta) - 1 - 2\frac{C}{\sqrt{N}} \\ &> 2\left(1 - \frac{\eta}{4}\right) - 1 - 2\frac{\eta}{4} = 1 - \eta \end{aligned}$$

One can easily choose $\alpha > \alpha_\eta$ to extend the inequality $P(-\alpha_\eta < G_N f_{t_0} \leq \alpha_\eta) > 1 - \eta$ to all $N \geq 1$. This finishes the proof of point 1.

The proof of point 2, is a consequence of assumption **(L)**. For all $t, s \in I$, one has the following inequality

$$\begin{aligned} & \mathbb{E} \left[|G_N f_t - G_N f_s|^2 \right] \\ &= \mathbb{E} \left[|(f_t(X) - P f_t) - (f_s(X) - P f_s)|^2 \right] \\ &\leq (2k|t - s|)^2 \\ &= (2kt - 2ks)^2. \end{aligned}$$

This proves point 2. with $\gamma = 2, \alpha = 2$ and $\psi(u) = 2ku$, and finishes the proof of Theorem 2. \square

The proof of Gaussian approximation is based on a result from [2]. They derive rates for the Gaussian approximation of more general processes. They do so by approaching the process by finite-dimensional marginals and applying a multidimensional Gaussian approximation result. Controlling the covering number of the family (or equivalently its metric entropy) allows them to derive a good trade-off between a good approximation by the marginals and keeping the dimension small enough to obtain good rates in the multidimensional Gaussian approximation. For the sake of completeness, let us recall their result.

Let \mathcal{M} be the set of all measurable real-valued functions on $(\mathbb{X}, \mathcal{X})$. The authors work with a centered and scaled empirical process $\{G_N \tilde{f}, \tilde{f} \in \tilde{\mathcal{F}}\}$ indexed by a general family $\tilde{\mathcal{F}} \subset \mathcal{M}$, not necessarily indexed by I . Their result shows that, under mild assumptions, this process can be approached by a centered Gaussian process $\tilde{\mathbb{G}}$ indexed by $\tilde{\mathcal{F}}$ with covariance $\mathbb{E}[\tilde{\mathbb{G}}(\tilde{f})\tilde{\mathbb{G}}(\tilde{g})] = P\tilde{f}\tilde{g} - P\tilde{f}P\tilde{g}$, for $\tilde{f}, \tilde{g} \in \tilde{\mathcal{F}}$. Let us define the covering number of $\tilde{\mathcal{F}}$. First, assume that there exists $\tilde{F} \in \mathcal{M}$ such that for all $\tilde{f} \in \tilde{\mathcal{F}}$ and all $x \in \mathbb{X}$, $|\tilde{f}(x)| \leq \tilde{F}(x)$. We say that \tilde{F} is an *envelope* of $\tilde{\mathcal{F}}$. For all probability measures Q on $(\mathbb{X}, \mathcal{X})$, we consider the semi-metric $d_Q(g, h)^2 = \int (g - h)^2 dQ$, for $g, h \in \mathcal{M}$. Under d_Q , we define the ball of radius $\delta > 0$ centered in $h \in \mathcal{M}$ by $B_Q(h, \delta) := \{g \in \mathcal{M}, d_Q(g, h) < \delta\}$. We also define $\tilde{F}_Q^2 := \int \tilde{F}^2 dQ$. For $\delta > 0$, let $N(\tilde{\mathcal{F}}, d_Q, \delta)$ be the size of the smallest finite subset $K \subset \mathcal{M}$ verifying that the union of balls $B_Q(h, \delta)$ for h in K covers $\tilde{\mathcal{F}}$. Finally, we set the covering number of $\tilde{\mathcal{F}}$ to be $N(\tilde{\mathcal{F}}, \delta) := \sup_Q N(\tilde{\mathcal{F}}, d_Q, \delta \cdot \tilde{F}_Q)$, where the supremum is taken over all Q such that $0 < \tilde{F}_Q < \infty$.

Before stating the result of [2], consider these two basic assumptions.

- (F.i) For some $M > 0$ and for all $\tilde{f} \in \tilde{\mathcal{F}}$, $\sup_{x \in \mathbb{X}} |\tilde{f}(x)| \leq M/2$.
- (F.ii) The class $\tilde{\mathcal{F}}$ is point-wise measurable, i.e., there exists a countable subclass $\tilde{\mathcal{F}}_\infty$ of $\tilde{\mathcal{F}}$ such that we can find for any function $\tilde{f} \in \tilde{\mathcal{F}}$ a sequence of functions $\{\tilde{f}_m\}$ in $\tilde{\mathcal{F}}_\infty$ for which $\lim_{m \rightarrow \infty} \tilde{f}_m(x) = \tilde{f}(x)$ for all $x \in \mathbb{X}$.

Proposition 6. [2, Proposition 1.] *Assume that $\tilde{\mathcal{F}}$ verifies (F.i) and (F.ii). Take $\tilde{F} := M/2$ as the envelope of $\tilde{\mathcal{F}}$. Moreover, assume that there exist positive constants c_0 and v_0 such that $N(\tilde{\mathcal{F}}, \delta) \leq c_0 \delta^{-v_0}$. Then, for each $\lambda > 1$ there is a constant $\rho(\lambda)$ such that for each N , one can construct X_1, \dots, X_N and $\tilde{\mathbb{G}}^{(N)}$ such that*

$$\mathbb{P} \left(\sup_{\tilde{f} \in \tilde{\mathcal{F}}} \left| G_N \tilde{f} - \tilde{\mathbb{G}}^{(N)}(\tilde{f}) \right| > \rho(\lambda) N^{-\frac{1}{2+5v_0}} \log N^{\frac{4+5v_0}{4+10v_0}} \right) \leq N^{-\lambda}.$$

Let us now prove the Gaussian approximation in Theorem 2 by applying the previous proposition.

Proof. The proof consists in applying Proposition 6 to \mathcal{F} with $v_0 = 1$. Clearly, assumption **(B)** implies (F.i). Since the paths $t \rightarrow f_t(x)$ are continuous for all $x \in \mathbb{X}$, taking $\mathcal{F}_\infty = \{f_t, t \in I \cap \mathbb{Q}\}$ gives (F.ii), where \mathbb{Q} is the set of rational numbers.

Let us prove the upper-bound on the covering number. Let L be the length of I . From **(L)**, we know that there exists a constant k , such that $t \rightarrow f_t(x)$ is k -lipschitz, for all x . Considering a regular grid on I , $\min I = t_0 < t_1 < \dots < t_q = \max I$. For all $t \in I$ there exists a integer j such that for all $x \in \mathbb{X}$, $|f_t(x) - f_{t_j}(x)| \leq k|t - t_j| \leq kL/q$. Hence, for all probability measures Q , $d_Q(f_t, f_{t_j}) \leq kL/q$, meaning that $N(\mathcal{F}, d_Q, kL/q) \leq q + 1$. As this last inequality stands for all Q , we can find a constant $c_0 > 0$ such that $N(\mathcal{F}, \delta) \leq c_0 \delta^{-1}$, for all $\delta > 0$. This finishes the proof. \square

B Application to the Heat distance processes.

B.1 Bounds for the laplacian eigenvalues.

Let us start by proving a lemma on the extremal laplacian eigenvalues of graphs in \mathcal{G}_n . We define Λ_{\min} and Λ_{\max} as, respectively, the minimal and maximal positive laplacian eigenvalues of graphs in \mathcal{G}_n :

$$\begin{aligned} \Lambda_{\min} &:= \inf\{\lambda > 0, \text{ s.t. } \lambda \text{ is an eigenvalue of } L(G), G \in \mathcal{G}_n\} \\ \Lambda_{\max} &:= \sup\{\lambda > 0, \text{ s.t. } \lambda \text{ is an eigenvalue of } L(G), G \in \mathcal{G}_n\}. \end{aligned}$$

Lemma 1. Λ_{\min} and Λ_{\max} satisfy the following bounds:

$$\Lambda_{\min} \geq \frac{8w_{\min}}{n^2}, \tag{B.1}$$

$$\Lambda_{\max} \leq nw_{\max}. \tag{B.2}$$

Note that (B.1) will not be used in the rest of the paper. Still, we choose to present it as we believe it could be of future use.

Proof. Take $G \in \mathcal{G}_n$, we want to prove that λ_{min} , the smallest positive eigenvalue of $L(G)$, verifies $\lambda_{min} \geq 8w_{min}n^{-2}$. To do so, we apply a Cheeger-type inequality. First assume that G is connected, hence $\lambda_{min} = \lambda_2(G)$. Let V be the set of vertices of G . Following [17, Section 3], we define the *average minimal cut* of G by

$$\gamma(G) = \min_{\emptyset \neq U \subsetneq V} \sum_{i \in U, j \in V \setminus U} \frac{w_{i,j}}{|U|(n-|U|)}.$$

As G is connected, there exists at least one edge with a weight greater than w_{min} joining U and $V \setminus U$. Hence,

$$\sum_{i \in U, j \in V \setminus U} w_{i,j} \geq w_{min}.$$

Moreover, for all U , $|U|(n-|U|) \leq n^2/4$. This yields to $\gamma(G) \geq 4w_{min}/n^2$. From [17, Theorem 2], we have that $\lambda_2(G) \geq 2\gamma(G)$, hence $\lambda_2(G) \geq 8w_{min}n^{-2}$. If now G is not connected, one can check that there exists G_{sub} a connected subgraph of G of size n_{sub} , such that $\lambda_{min} = \lambda_2(G_{sub})$. This gives

$$\lambda_{min} = \lambda_2(G_{sub}) \geq \frac{8w_{min}}{n_{sub}^2} \geq \frac{8w_{min}}{n^2}.$$

This finishes the proof of the first bound.

We now prove the second bound. The largest eigenvalue of $L(G)$ is denoted by λ_n . Letting x be an eigenvector associated to λ_n verifying $\|x\|_2 = 1$, we have

$$\lambda_n = x^T L(G)x = \sum_{i < j} w_{i,j}(x_i - x_j)^2 \leq w_{max} \sum_{i < j} (x_i - x_j)^2 = w_{max} x^T L(G_{comp})x$$

where G_{comp} is the complete undirected graph, hence

$$L(G_{comp}) = \begin{pmatrix} n-1 & -1 & \cdots & -1 \\ -1 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & -1 \\ -1 & \cdots & -1 & n-1 \end{pmatrix}.$$

One can check that $x^T L(G_{comp})x \leq \lambda_n(L(G_{comp})) = n$, so $\lambda_n \leq nw_{max}$. \square

B.2 Result on HKD processes.

To apply Theorem 2 to the HKD processes we first prove that they are uniformly bounded and Lipschitz-continuous.

Proposition 7. *For all G, G' in \mathcal{G}_n , denote by $(\lambda_k)_{1 \leq k \leq n}$ and $(\phi_k)_{1 \leq k \leq n}$ (resp. $(\lambda'_l)_{1 \leq l \leq n}$ and $(\phi'_l)_{1 \leq l \leq n}$) the eigenvalues and orthonormal eigenvectors of $L(G)$ (resp. $L(G')$). Remember that we always choose ϕ_1 and ϕ'_1 equal to the vector with all entries equal to $1/\sqrt{n}$. We denote by $\langle \cdot, \cdot \rangle$ the standard scalar product in \mathbb{R}^n . Then, the application $t \rightarrow D_t((G, G'))$ verifies:*

1. For all $t \in [0, T]$, $D_t((G, G'))$ can be written in terms of the eigen-elements of $L(G)$ and $L(G')$:

$$D_t((G, G')) = \left(\sum_{k,l=2}^n (e^{-t\lambda_k} - e^{-t\lambda'_l})^2 \langle \phi_k, \phi'_l \rangle^2 \right)^{1/2}. \quad (\text{B.3})$$

2. For all $t \in [0, T]$, $D_t((G, G')) \leq \sqrt{n}$.
3. $t \mapsto D_t((G, G'))$ is $(n^{3/2}w_{max})$ -Lipschitz continuous on $[0, T]$.

Proof. Let G, G' be in \mathcal{G}_n . We start by proving (B.3). This is done through the following computation:

$$\begin{aligned} & \|e^{-tL} - e^{-tL'}\|_F^2 \\ &= \|e^{-tL}\|_F^2 + \|e^{-tL'}\|_F^2 - 2\text{Tr}\left(e^{-tL}e^{-tL'}\right) \\ &= \sum_{k=1}^n e^{-2t\lambda_k} \|\phi_k \phi_k^T\|_F^2 + \sum_{l=1}^n e^{-2t\lambda'_l} \|\phi'_l \phi'_l{}^T\|_F^2 - 2 \sum_{k,l=1}^n e^{-t\lambda_k} e^{-t\lambda'_l} \text{Tr}\left(\phi_k \phi_k^T \phi'_l \phi'_l{}^T\right) \end{aligned}$$

One can prove that $\text{Tr}\left(\phi_k \phi_k^T \phi'_l \phi'_l{}^T\right) = \langle \phi_k, \phi'_l \rangle^2$.

Similarly $\|\phi_k \phi_k^T\|_F^2 = \|\phi_k\|_2^2 = \sum_{l=1}^n \langle \phi_k, \phi'_l \rangle^2$ and $\|\phi'_l \phi'_l{}^T\|_F^2 = \|\phi'_l\|_2^2 = \sum_{k=1}^n \langle \phi_k, \phi'_l \rangle^2$. So

$$\|e^{-tL} - e^{-tL'}\|_F^2 = \sum_{k,l=1}^n (e^{-t\lambda_k} - e^{-t\lambda'_l})^2 \langle \phi_k, \phi'_l \rangle^2.$$

The sums can start at $k = 2$ and $l = 2$ thanks to the facts that $\lambda_1 = \lambda'_1 = 0$ and $\phi_1 = \phi'_1$ combined with the orthogonality of the eigenvectors families. This finishes the proof of (B.3).

Bounding all terms $(e^{-t\lambda_k} - e^{-t\lambda'_l})^2$ by 1 in (B.3), and using the orthonormality of the eigenvectors families yields to $D_t((G, G')) \leq \sqrt{n}$.

We now prove the Lipschitz result. One can check that $t \rightarrow D_t((G, G'))$ is \mathcal{C}^1 on $[0, T]$. The case $G = G'$ is easily dealt with. Assume now that $G \neq G'$. For all $t \in (0, T]$,

$$\begin{aligned} \left| \frac{d}{dt} D_t((G, G')) \right| &= \left| \frac{- \sum_{k,l=2}^n (e^{-t\lambda_k} - e^{-t\lambda'_l}) (\lambda_k e^{-t\lambda_k} - \lambda'_l e^{-t\lambda'_l}) \langle \phi_k, \phi'_l \rangle^2}{\left(\sum_{k,l=2}^n (e^{-t\lambda_k} - e^{-t\lambda'_l})^2 \langle \phi_k, \phi'_l \rangle^2 \right)^{1/2}} \right| \\ &\leq \sum_{k,l=2}^n \frac{|e^{-t\lambda_k} - e^{-t\lambda'_l}| |\lambda_k e^{-t\lambda_k} - \lambda'_l e^{-t\lambda'_l}| \langle \phi_k, \phi'_l \rangle^2}{\left(\sum_{i,j=2}^n (e^{-t\lambda_i} - e^{-t\lambda'_j})^2 \langle \phi_i, \phi'_j \rangle^2 \right)^{1/2}} \\ &\leq \sqrt{\sum_{k,l=2}^n |\lambda_k e^{-t\lambda_k} - \lambda'_l e^{-t\lambda'_l}|^2 \langle \phi_k, \phi'_l \rangle^2} \sqrt{\frac{\sum_{k,l=2}^n (e^{-t\lambda_k} - e^{-t\lambda'_l})^2 \langle \phi_k, \phi'_l \rangle^2}{\sum_{i,j=2}^n (e^{-t\lambda_i} - e^{-t\lambda'_j})^2 \langle \phi_i, \phi'_j \rangle^2}} \\ &= \sqrt{\sum_{k,l=2}^n |\lambda_k e^{-t\lambda_k} - \lambda'_l e^{-t\lambda'_l}|^2 \langle \phi_k, \phi'_l \rangle^2}, \end{aligned}$$

where the last inequality comes from the Cauchy-Schwarz inequality.

According to the mean value theorem, for all k and l

$$|\lambda_k e^{-t\lambda_k} - \lambda'_l e^{-t\lambda'_l}| \leq |\lambda_k - \lambda'_l| \leq \Lambda_{\max}.$$

Hence, for all $t \in (0, T]$,

$$\left| \frac{d}{dt} D_t((G, G')) \right| \leq \Lambda_{\max} \sqrt{\sum_{k,l=2}^n \langle \phi_k, \phi'_l \rangle^2} \leq n w_{\max} \sqrt{n} = n^{3/2} w_{\max}.$$

This finishes the proof. \square

The proof of Theorem 4 is a direct application of Theorem 2, as Proposition 7 ensures that the hypotheses of Theorem 2 are verified by the HKD process.

B.3 Result on HPD processes.

To prove Theorem 5, we start by showing that under the hypotheses of the theorem the HPD processes are uniformly bounded and Lipschitz-continuous.

Proposition 8. *For all G, G' in \mathcal{G}^n , the application $t \rightarrow H_t((G, G'))$ verifies:*

1. *For all $t \in [0, T]$, $0 \leq H_t((G, G')) \leq 1$.*
2. *$t \mapsto H_t((G, G'))$ is $(2nw_{\max})$ -Lipschitz-continuous on $[0, T]$.*

Proof. Recall that for all $G \in \mathcal{G}^n$, and for a vertex i ,

$$h_t(G)(i) = \sum_{k=1}^{n(G)} e^{-t\lambda_k} \phi_k(i)^2,$$

where $(\lambda_k)_{1 \leq k \leq n(G)}$ and $(\phi_k)_{1 \leq k \leq n(G)}$ are the eigenvalues and orthonormal eigenvectors of $L(G)$ and $n(G)$ is the size of G . Hence $0 \leq h_t(G)(i) \leq 1$ for all i , meaning that all points in the diagram $Dg(G, h_t(G))$ are contained in $[0, 1]^2$. So from the definition of the Bottleneck distance, $0 \leq H_t((G, G')) \leq 1$, for all G, G' in \mathcal{G}^n and for all $t \in [0, T]$.

Let us now compute the first derivative of $h_t(G)(i)$:

$$\frac{d}{dt} h_t(G)(i) = - \sum_{k=1}^{n(G)} \lambda_k e^{-t\lambda_k} \phi_k(i)^2.$$

Its absolute value is upper-bounded by $\lambda_{n(G)}$, the largest eigenvalue of $L(G)$. From Lemma 1, we have $\lambda_{n(G)} \leq n(G)w_{\max}$. Hence, $t \rightarrow h_t(G)$ is (nw_{\max}) -Lipschitz continuous on $[0, T]$. To conclude, we come back to the definition of the HPD in terms of distance between persistence diagrams. Applying the triangular inequality to the Bottleneck distance gives for all $G, G' \in \mathcal{G}^n$, for all $t, t' \in [0, T]$, and for all diagram construction Dg ,

$$\begin{aligned} & |d_B(Dg(G, h_t(G)), Dg(G', h_t(G'))) - d_B(Dg(G, h_{t'}(G)), Dg(G', h_{t'}(G')))| \\ & \leq d_B(Dg(G, h_t(G)), Dg(G, h_{t'}(G))) + d_B(Dg(G', h_t(G')), Dg(G', h_{t'}(G'))). \end{aligned}$$

Similarly, the same inequality but with maxima over the diagram constructions holds:

$$\begin{aligned} & |H_t((G, G')) - H_{t'}((G, G'))| \\ & \leq \max d_B(Dg(G, h_t(G)), Dg(G, h_{t'}(G))) + \max d_B(Dg(G', h_t(G')), Dg(G', h_{t'}(G'))). \end{aligned}$$

Applying Theorem 1 and using the Lipschitz continuity of the HKS yields

$$\begin{aligned} & |H_t((G, G')) - H_{t'}((G, G'))| \\ & \leq \|h_t(G) - h_{t'}(G)\|_\infty + \|h_t(G') - h_{t'}(G')\|_\infty \\ & \leq 2nw_{\max} |t - t'|. \end{aligned}$$

□

Theorem 5 follows from Proposition 8 and Theorem 2.

Acknowledgments

I am thankful to Frédéric Chazal¹ and Pascal Massart² for the valuable discussions and advice on this work. I also want to thank the people from Datashape¹ and LMO² for the enriching conversations.

¹Inria Saclay

²Université Paris-Saclay

References

- [1] ALI, W., RITO, T., REINERT, G., SUN, F., AND DEANE, C. M. Alignment-free protein interaction network comparison. *Bioinformatics* 30, 17 (2014), i430–i437.
- [2] BERTHET, P., AND MASON, D. M. Revisiting two strong approximation results of dudley and philipp. In *High dimensional probability*. Institute of Mathematical Statistics, 2006, pp. 155–172.
- [3] BILLINGSLEY, P. *Convergence of probability measures*. John Wiley & Sons, 2013.
- [4] BORGWARDT, K. M., AND KRIEGEL, H.-P. Shortest-path kernels on graphs. In *Fifth IEEE international conference on data mining (ICDM'05)* (2005), IEEE, pp. 8–pp.
- [5] CARRIÈRE, M., CHAZAL, F., IKE, Y., LACOMBE, T., ROYER, M., AND UMEDA, Y. Perslay: a neural network layer for persistence diagrams and new graph topological signatures. In *International Conference on Artificial Intelligence and Statistics* (2020), PMLR, pp. 2786–2796.
- [6] CASTELLAZZI, G., DEBERNARD, L., MELZER, T. R., DALRYMPLE-ALFORD, J. C., D'ANGELO, E., MILLER, D. H., GANDINI WHEELER-KINGSHOTT, C. A., AND MASON, D. F. Functional connectivity alterations reveal complex mechanisms based on clinical and radiological status in mild relapsing remitting multiple sclerosis. *Frontiers in neurology* 9 (2018), 690.
- [7] CHAZAL, F., DE SILVA, V., GLISSE, M., AND OUDOT, S. *The structure and stability of persistence modules*. Springer, 2016.
- [8] COHEN-STEINER, D., EDELSBRUNNER, H., AND HARER, J. Extending persistence using poincaré and lefschetz duality. *Foundations of Computational Mathematics* 9, 1 (2009), 79–103.
- [9] COIFMAN, R. R., AND HIRN, M. J. Diffusion maps for changing data. *Applied and computational harmonic analysis* 36, 1 (2014), 79–107.
- [10] COIFMAN, R. R., AND LAFON, S. Diffusion maps. *Applied and computational harmonic analysis* 21, 1 (2006), 5–30.
- [11] EDELSBRUNNER, H., AND HARER, J. *Computational topology: an introduction*. American Mathematical Soc., 2010.
- [12] EMMERT-STREIB, F., DEHMER, M., AND SHI, Y. Fifty years of graph matching, network alignment and network comparison. *Information sciences* 346 (2016), 180–197.
- [13] ERDOS, P., RÉNYI, A., ET AL. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci* 5, 1 (1960), 17–60.
- [14] FAISAL, F. E., NEWAZ, K., CHANEY, J. L., LI, J., EMRICH, S. J., CLARK, P. L., AND MILENKOVIĆ, T. Grafene: Graphlet-based alignment-free network approach integrates 3d structural and sequence (residue order) data to improve protein structural comparison. *Scientific reports* 7, 1 (2017), 1–15.
- [15] FAIVRE, A., ROBINET, E., GUYE, M., ROUSSEAU, C., MAAROUF, A., LE TROTTER, A., ZAARAOU, W., RICO, A., CRESPIY, L., SOULIER, E., ET AL. Depletion of brain functional connectivity enhancement leads to disability progression in multiple sclerosis: a longitudinal resting-state fmri study. *Multiple Sclerosis Journal* 22, 13 (2016), 1695–1708.
- [16] FARAHANI, F. V., KARWOWSKI, W., AND LIGHTHALL, N. R. Application of graph theory for identifying connectivity patterns in human brain networks: a systematic review. *frontiers in Neuroscience* 13 (2019), 585.
- [17] FIEDLER, M. An estimate for the nonstochastic eigenvalues of doubly stochastic matrices. *Linear algebra and its applications* 214 (1995), 133–143.
- [18] GERA, R., ALONSO, L., CRAWFORD, B., HOUSE, J., MENDEZ-BERMUDEZ, J., KNUTH, T., AND MILLER, R. Identifying network structure similarity using spectral graph theory. *Applied network science* 3, 1 (2018), 1–15.
- [19] GRETTON, A., BORGWARDT, K. M., RASCH, M. J., SCHÖLKOPF, B., AND SMOLA, A. A kernel two-sample test. *The Journal of Machine Learning Research* 13, 1 (2012), 723–773.
- [20] HAMMOND, D. K., GUR, Y., AND JOHNSON, C. R. Graph diffusion distance: A difference measure for weighted graphs based on the graph laplacian exponential kernel. In *2013 IEEE Global Conference on Signal and Information Processing* (2013), IEEE, pp. 419–422.
- [21] HIDO, S., AND KASHIMA, H. A linear-time graph kernel. In *2009 Ninth IEEE International Conference on Data Mining* (2009), IEEE, pp. 179–188.

- [22] HOLLAND, P. W., LASKEY, K. B., AND LEINHARDT, S. Stochastic blockmodels: First steps. *Social networks* 5, 2 (1983), 109–137.
- [23] HU, N., RUSTAMOV, R. M., AND GUIBAS, L. Stable and informative spectral signatures for graph matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2014), pp. 2305–2312.
- [24] JOHANSSON, F., JETHAVA, V., DUBHASHI, D., AND BHATTACHARYYA, C. Global graph kernels using geometric embeddings. In *International Conference on Machine Learning* (2014), PMLR, pp. 694–702.
- [25] KOSOROK, M. R. Introduction to empirical processes. *Introduction to Empirical Processes and Semiparametric Inference* (2008).
- [26] KOUTRA, D., VOGELSTEIN, J. T., AND FALOUTSOS, C. Deltacon: A principled massive-graph similarity function. In *Proceedings of the 2013 SIAM International Conference on Data Mining* (2013), SIAM, pp. 162–170.
- [27] KRIEGE, N. M., JOHANSSON, F. D., AND MORRIS, C. A survey on graph kernels. *Applied Network Science* 5, 1 (2020), 1–42.
- [28] LEHMANN, E. L., AND ROMANO, J. P. *Testing statistical hypotheses*. Springer Science & Business Media, 2006.
- [29] MARCOTTE, S., BARBE, A., GRIBONVAL, R., VAYER, T., SEBBAN, M., BORGNAT, P., AND GONÇALVES, P. Fast multiscale diffusion on graphs. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2022), IEEE, pp. 5627–5631.
- [30] MARIA, C., BOISSONNAT, J.-D., GLISSE, M., AND YVINEC, M. The gudhi library: Simplicial complexes and persistent homology. In *International congress on mathematical software* (2014), Springer, pp. 167–174.
- [31] NAKATSUKASA, Y., AND HIGHAM, N. J. Stable and efficient spectral divide and conquer algorithms for the symmetric eigenvalue decomposition and the svd. *SIAM Journal on Scientific Computing* 35, 3 (2013), A1325–A1349.
- [32] NEYMAN, J., AND PEARSON, E. S. IX. on the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character* 231, 694–706 (1933), 289–337.
- [33] OUDOT, S. Y. *Persistence theory: from quiver representations to data analysis*, vol. 209. American Mathematical Society Providence, 2015.
- [34] PENROSE, M., ET AL. *Random geometric graphs*, vol. 5. Oxford university press, 2003.
- [35] PRŽULJ, N. Biological network comparison using graphlet degree distribution. *Bioinformatics* 23, 2 (2007), e177–e183.
- [36] PRŽULJ, N., CORNEIL, D. G., AND JURISICA, I. Modeling interactome: scale-free or geometric? *Bioinformatics* 20, 18 (2004), 3508–3515.
- [37] ROCCA, M. A., VALSASINA, P., MEANI, A., FALINI, A., COMI, G., AND FILIPPI, M. Impaired functional integration in multiple sclerosis: a graph theory study. *Brain Structure and Function* 221, 1 (2016), 115–131.
- [38] SHERVASHIDZE, N., SCHWEITZER, P., VAN LEEUWEN, E. J., MEHLHORN, K., AND BORGWARDT, K. M. Weisfeiler-lehman graph kernels. *Journal of Machine Learning Research* 12, 9 (2011).
- [39] SHERVASHIDZE, N., VISHWANATHAN, S., PETRI, T., MEHLHORN, K., AND BORGWARDT, K. Efficient graphlet kernels for large graph comparison. In *Artificial intelligence and statistics* (2009), PMLR, pp. 488–495.
- [40] SIGLIDIS, G., NIKOLENTZOS, G., LIMNIOS, S., GIATSIDIS, C., SKIANIS, K., AND VAZIRGIANNIS, M. Grakel: A graph kernel library in python. *The Journal of Machine Learning Research* 21, 1 (2020), 1993–1997.
- [41] SOUNDARAJAN, S., ELIASSI-RAD, T., AND GALLAGHER, B. A guide to selecting a network similarity method. In *Proceedings of the 2014 Siam international conference on data mining* (2014), SIAM, pp. 1037–1045.
- [42] SUGIYAMA, M., AND BORGWARDT, K. Halting in random walk kernels. *Advances in neural information processing systems* 28 (2015).

- [43] SUN, J., OVSJANIKOV, M., AND GUIBAS, L. A concise and provably informative multi-scale signature based on heat diffusion. In *Computer graphics forum* (2009), vol. 28, Wiley Online Library, pp. 1383–1392.
- [44] TANTARDINI, M., IEVA, F., TAJOLI, L., AND PICCARDI, C. Comparing methods for comparing networks. *Scientific reports* 9, 1 (2019), 1–19.
- [45] TSITSULIN, A., MOTTIN, D., KARRAS, P., BRONSTEIN, A., AND MÜLLER, E. Netlsd: hearing the shape of a graph. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (2018), pp. 2347–2356.
- [46] VAN DER VAART, A. W., AND WELLNER, J. A. Weak convergence. In *Weak convergence and empirical processes*. Springer, 1996.
- [47] WATTS, D. J., AND STROGATZ, S. H. Collective dynamics of ‘small-world’ networks. *nature* 393, 6684 (1998), 440–442.
- [48] WILSON, R. C., AND ZHU, P. A study of graph spectra for comparing graphs and trees. *Pattern Recognition* 41, 9 (2008), 2833–2841.
- [49] YAVEROĞLU, Ö. N., MALOD-DOGNIN, N., DAVIS, D., LEVNAJIC, Z., JANJIC, V., KARAPANDZA, R., STOJMIROVIC, A., AND PRŽULJ, N. Revealing the hidden language of complex networks. *Scientific reports* 4, 1 (2014), 1–9.