

# Bridging Distance and Spectral Positional Encodings via Anchor-Based Diffusion Geometry Approximation

Anonymous ACL submission

## Abstract

Molecular graph learning benefits from positional signals that capture both local neighborhoods and global topology. Two widely used families are spectral encodings derived from Laplacian or diffusion operators and anchor-based distance encodings built from shortest-path information, yet their precise relationship is poorly understood. We interpret distance encodings as a low-rank surrogate of diffusion geometry and derive an explicit trilateration map that reconstructs truncated diffusion coordinates from transformed anchor distances and anchor spectral positions, with pointwise and Frobenius-gap guarantees on random regular graphs. On DrugBank molecular graphs using a shared GNP-based DDI prediction backbone, a distance-driven Nyström scheme closely recovers diffusion geometry, and both Laplacian and distance encodings substantially outperform a no-encoding baseline.

## 1 Introduction

Molecular graph modeling tasks—from de novo graph generation to molecular reasoning—require representations that capture both local chemical neighborhoods and global topology (2017; 2020; 2020). A recurring challenge is how to inject structural positional information so that long-range dependencies and global organization are accessible to the model. In this work we use Graph Neural Processes (GNPs) as a flexible probabilistic backbone (2025c).

Many positional/structural signals have been explored, including random-walk/propagation statistics, structural-role descriptors, and relative-position biases in graph transformers (2022; 2023; 2021). We focus on two particularly common families. Spectral/diffusion encodings (e.g., Laplacian eigenmaps, heat-kernel and diffusion-map features) provide coordinates aligned with diffusion geometry but often rely on costly eigendecomposition

or careful approximation (2007; 2023; 2022). Distance/anchor encodings represent each node by its shortest-path distances to a small anchor set (optionally transformed), offering a simple and scalable alternative without explicit spectral computation (2025a; 2009).

Despite their widespread use, an explicit algebraic account of when anchor-based distance features can approximate diffusion geometry is still missing, as are error measures that are meaningful both for recovered coordinates and for induced distance/kernel matrices. This gap matters in molecular settings where diffusion geometry is a natural inductive bias and computational budgets are limited, raising the practical question of when distance encodings can serve as principled substitutes for spectral encodings in tasks such as drug-drug interaction (DDI) prediction (2018; 2018).

**Motivation.** We ask: *Can we construct an explicit algebraic map from anchor-based shortest-path encodings to truncated diffusion (spectral) coordinates, with provable pointwise and matrix-level error guarantees?* Such a link would connect discrete distance primitives to diffusion geometry and clarify the accuracy-efficiency trade-off behind replacing spectral computation with distance features.

We address this question through theory and experiments. On the theory side, we derive an explicit trilateration operator that reconstructs truncated diffusion coordinates from transformed anchor distances and anchor spectral positions, and we establish pointwise and Frobenius-gap guarantees under a random regular graph model with a local monotone distance linkage. On the empirical side, we evaluate diffusion-geometry recovery on DrugBank using a distance-driven Nyström scheme and assess downstream DDI prediction on DrugBank and ChCh-Miner under a shared GNP backbone (2020; 2020; 2023). Overall, both Laplacian and distance encodings improve over NoPE, with Laplacian en-

083 codings giving the most consistent gains.

084 Our main contributions are:

- 085 1. We provide a spectral–algebraic bridge by formalizing anchor-based distance encodings as a low-rank surrogate for diffusion geometry and developing an explicit trilateration operator with pointwise and Frobenius-gap guarantees under random regular graph assumptions.
- 086 087 088 089 090 091 092 093 094 095 2. We show that distance-based Nyström approximations closely recover diffusion geometry on DrugBank molecular graphs, yielding accurate kernel and embedding approximations with a moderate number of anchors.
- 096 097 098 099 100 3. Using a common GNP-based architecture for DDI prediction, we compare NoPE, DE, and LapPE, and conduct ablations on distance transforms and anchor counts on DrugBank and ChCh-Miner.

## 101 2 Literature Review

102 This section reviews structural and positional information for graph learning, with an emphasis on molecular graphs. We discuss why purely local computation can be structurally limiting, summarize spectral/diffusion and anchor-based distance encodings as two common positional encoding families, and review drug-drug interaction (DDI) models to identify the gap addressed by our spectral-algebraic bridge between distance features and diffusion geometry.

### 112 2.1 Graph structure learning and locality limitations

113 114 Many graph learning methods adopt locality-based computation, most prominently message passing: node representations are iteratively updated by aggregating and transforming features from adjacent nodes, forming the backbone of many GNN architectures and related graph encoders (2025c; 2017; 2016). Such designs have shown strong empirical performance across a wide range of applications, including molecular property prediction.

123 124 125 126 127 128 129 130 Despite this success, locality imposes structural limitations. The expressive power of standard message passing is upper bounded by the Weisfeiler-Leman test, so non-isomorphic graphs and structurally distinct nodes can remain indistinguishable (2019; 2018). In addition, long-range dependencies can be difficult to capture due to over-smoothing and over-squashing effects (2021; 2022). These

issues motivate enriching node features with explicit structural or positional signals beyond local neighborhood aggregation.

### 2.2 Spectral positional encodings

Beyond spectral and distance-based designs, prior work also studies random-walk/propagation statistics, structural roles, transformer relative-position biases, and substructure-count descriptors (2022; 2023; 2021). We focus on spectral/diffusion and anchor-distance encodings because they provide a natural geometric lens and are the two families we explicitly connect.

Spectral positional encodings derive coordinates from Laplacian or diffusion-operator eigenvectors (e.g., Laplacian eigenmaps and diffusion maps) (2003; 2005; 2006), and are widely used as fixed Laplacian positional features to inject global topology and break symmetries (2023). Variants adapt or approximate these features via learned reweighting/combinations or propagation-based approximations (2022; 2023), but spectral methods can still be costly on large or evolving graphs due to reliance on global modes.

### 2.3 Distance based positional encodings

Distance-based positional encodings represent each node by its shortest-path distances (or simple transforms) to a small set of anchor nodes (2009; 2019). By relying on shortest-path computations rather than eigen-decomposition, these features are typically easy to implement and scale well with standard graph routines.

Such encodings have been used to enhance expressivity beyond Weisfeiler-Leman (2009) and to build structure-aware graph transformers through shortest-path distances and related structural priors (2021). However, their connection to spectral/diffusion geometry is often only discussed qualitatively, leaving open when a finite anchor set can approximate diffusion geometry and how to quantify approximation error at both coordinate and kernel/distance levels.

### 2.4 Graph based models for drug-drug interaction prediction

Graph-based approaches have become central to drug-drug interaction (DDI) prediction. Early deep learning methods combined molecular fingerprints or sequence representations with standard neural architectures (2018), while later work leveraged biomedical knowledge graphs and applied GNNs

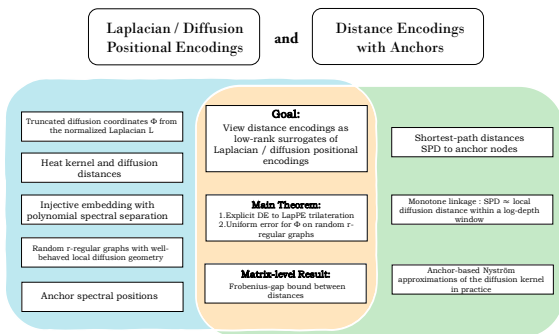


Figure 1: Overview of spectral/diffusion and anchor-distance positional encodings, and the algebraic bridge developed in this work.

or knowledge-graph neural networks to model interactions in heterogeneous networks (2018; 2020). Multimodal frameworks further enriched drug representations by integrating molecular graphs with targets, pathways, and text (2020).

More recent architectures explicitly exploit molecular structure via dual-graph designs and co-attention mechanisms to capture substructure-level interactions between drug pairs (2025b; 2025). These results consistently highlight the importance of structural information for accurate DDI prediction; nevertheless, most models adopt a fixed positional/structural feature design and emphasize architectural or data-integration choices, with comparatively limited work that systematically relates distance-based encodings to spectral diffusion geometry under a unified theoretical lens.

## 2.5 Summary and remaining gap

Spectral encodings yield geometrically meaningful but globally costly coordinates, whereas distance-based encodings are scalable and simple yet are typically justified empirically and lack a unified algebraic account of when they serve as low-rank surrogates for diffusion geometry. We close this gap by providing reconstruction operators and error bounds linking distance and spectral encodings and evaluating the resulting accuracy-efficiency trade-off for DDI prediction in a unified GNP-based framework (Figures 1 and 2).

## 3 Preliminaries

We collect notation and basic geometric objects used in the analysis, and formalize the two positional-encoding families, anchor-based shortest-path distance encodings and truncated Laplacian spectral coordinates, together with the

algebraic comparison problem studied in Section 4.

### 3.1 Notations

Let  $G = (V, E)$  be a finite, connected, undirected graph with  $|V| = n$ . For  $u, v \in V$ , let  $\text{SPD}(u, v)$  denote the shortest-path distance. For  $u \in V$  and  $R \in \mathbb{N}$ , define the radius- $R$  ball

$$B_G(u, R) = \{v \in V : \text{SPD}(u, v) \leq R\}. \quad (1)$$

Let  $A \in \{0, 1\}^{n \times n}$  be the adjacency matrix and  $D$  the diagonal degree matrix. In the theoretical part we focus on  $r$ -regular graphs, so  $D = rI$  and we use

$$L := I - \frac{1}{r}A. \quad (2)$$

Let  $\{(\lambda_j, \varphi_j)\}_{j=1}^n$  be the eigenpairs of  $L$ , with  $\{\varphi_j\}$  orthonormal in  $\ell^2(V)$  and  $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ . For  $t > 0$ , define the heat semigroup

$$K_t = e^{-tL}, \quad (3)$$

with entries  $k_t(u, v) = (K_t)_{uv}$ . The diffusion distance at time  $t$  is equivalently given by

$$\begin{aligned} d_t(u, v)^2 &= \sum_{j=1}^n e^{-2t\lambda_j} (\varphi_j(u) - \varphi_j(v))^2 \\ &= k_{2t}(u, u) + k_{2t}(v, v) - 2k_{2t}(u, v). \end{aligned} \quad (4)$$

For an integer  $m \geq 1$ , define the truncated diffusion-map embedding

$$\Phi_t^{(m)}(v) = (e^{-t\lambda_{j+1}} \varphi_{j+1}(v))_{j=1}^m \in \mathbb{R}^m, \quad (5)$$

and the corresponding truncated diffusion distance

$$d_t^{(m)}(u, v) = \|\Phi_t^{(m)}(u) - \Phi_t^{(m)}(v)\|_2. \quad (6)$$

Given anchors  $a_1, \dots, a_{m+1} \in V$ , write

$$p_i = \Phi_t^{(m)}(a_i) \in \mathbb{R}^m, \quad i = 1, \dots, m+1. \quad (7)$$

### 3.2 Problem Statement

Fix a graph  $G = (V, E)$  and an anchor set  $S = \{a_1, \dots, a_{m+1}\} \subset V$ . For any node  $v \in V$ , define the (transformed) distance encoding

$$\tilde{\zeta}(v | S) = (\psi(\text{SPD}(a_1, v)), \dots, \psi(\text{SPD}(a_{m+1}, v)))^\top \in \mathbb{R}^{m+1}, \quad (8)$$

where  $\psi : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$  is a monotone scalar transform specified in Section 4; the raw distance encoding corresponds to  $\psi = \text{id}$ . The target Laplacian spectral encoding is the truncated coordinate

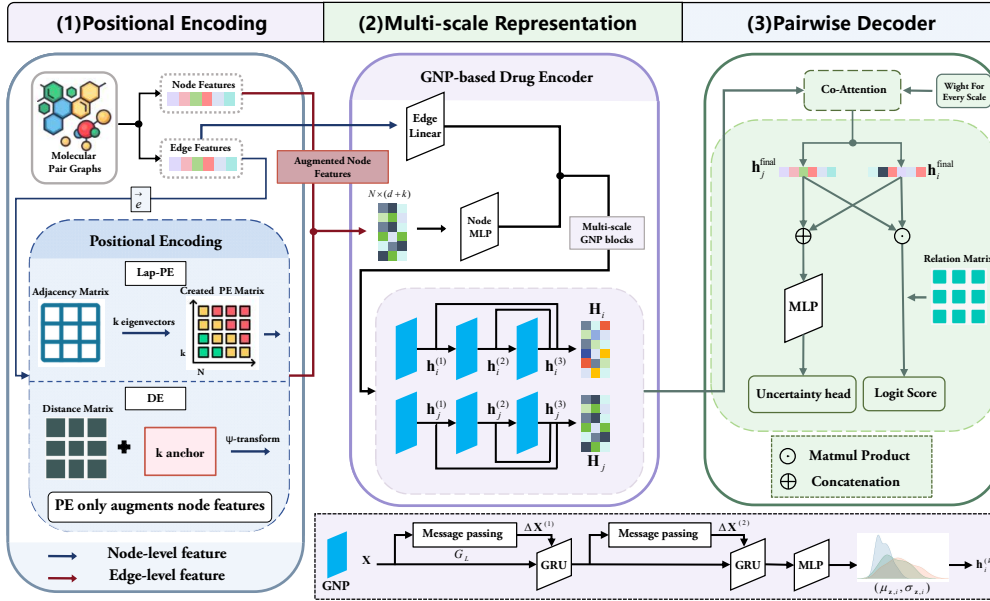


Figure 2: Proposed DDI framework: positional encodings augment molecular graphs, which are encoded by a shared multi-scale GNP and decoded via co-attention and relation-aware scoring.

$\Phi_t^{(m)}(v) \in \mathbb{R}^m$  (or a sign-/basis-invariant variant thereof).

Our central question is:

Under suitable structural conditions on  $G$  and a well-posed anchor configuration  $S$ , can one construct an explicit algebraic map that takes  $\zeta(v | S)$  as input and outputs  $\Phi_t^{(m)}(v)$  with quantitatively controlled error, and get a bound associated matrix-level discrepancy (e.g., in Frobenius norm) between the induced distance/kernel quantities?

In Section 4 we formalize these conditions and show that such a linkage exists, with error controlled (up to explicit constants) by a structural mismatch parameter  $\delta$ , in a logarithmic neighborhood regime consistent with common locality scales in prior analyses.

## 4 Theoretical Analysis

This section analyzes the algebraic relation between shortest-path distance encodings (DE) and truncated Laplacian spectral coordinates (LapPE) on  $r$ -regular graphs. All proofs are deferred to Appendix A.4. Throughout this section we use the (symmetric normalized) Laplacian in the  $r$ -regular form  $L = I - \frac{1}{r}A$ , so that the diffusion operator is  $K_t = e^{-tL}$  and the truncated diffusion-map embedding is  $\Phi_t^{(m)}$  as defined in Section 3.

### 4.1 Assumptions

**Assumption 1.** For the theoretical analysis, we consider graphs drawn from the random regular model  $G \sim \mathcal{G}_{n,r}$ , the uniform distribution on labeled  $r$ -regular graphs on  $[n]$  with fixed degree  $r \geq 3$ .

**Assumption 2.** There exist  $m \in \mathbb{N}$  and  $\alpha > 0$  such that, with high probability over  $G \sim \mathcal{G}_{n,r}$ , the truncated embedding  $\Phi_t^{(m)} : V \rightarrow \mathbb{R}^m$  is injective and satisfies

$$\min_{u \neq v} \|\Phi_t^{(m)}(u) - \Phi_t^{(m)}(v)\|_2 \geq n^{-\alpha}. \quad (9)$$

These assumptions hold with high probability for  $G \sim \mathcal{G}_{n,r}$  in the same logarithmic regime commonly used in analyses of Laplacian and distance-based positional encodings.

### 4.2 Main Theoretical Results

We first define an explicit trilateration operator that maps distance features to an  $m$ -dimensional spectral coordinate.

**Definition 1.** Fix  $t > 0$  and  $m \in \mathbb{N}$ , and let  $a_1, \dots, a_{m+1} \in V$  be anchors with

$$p_i = \Phi_t^{(m)}(a_i) \in \mathbb{R}^m, \quad i = 1, \dots, m+1. \quad (10)$$

Define

$$A = 2 \begin{pmatrix} (p_1 - p_{m+1})^\top \\ \vdots \\ (p_m - p_{m+1})^\top \end{pmatrix} \in \mathbb{R}^{m \times m}. \quad (11)$$

For any  $r = (r_1, \dots, r_{m+1})^\top \in \mathbb{R}^{m+1}$ , define

$$b(r) = \begin{pmatrix} \|p_1\|_2^2 - \|p_{m+1}\|_2^2 + r_{m+1}^2 - r_1^2 \\ \vdots \\ \|p_m\|_2^2 - \|p_{m+1}\|_2^2 + r_{m+1}^2 - r_m^2 \end{pmatrix} \in \mathbb{R}^m. \quad (12)$$

Let  $\psi : [0, R] \rightarrow \mathbb{R}_+$  be the strictly increasing scalar function from Theorem 2 and write  $\psi_*(y)$  for its elementwise application to  $y \in \mathbb{R}^{m+1}$ , i.e.,

$$(\psi_*(y))_i := \psi(y_i), \quad i = 1, \dots, m+1. \quad (13)$$

Recall the (raw) distance encoding to anchors  $S = \{a_1, \dots, a_{m+1}\}$ :

$$\zeta(v | S) = (\text{SPD}(a_1, v), \dots, \text{SPD}(a_{m+1}, v))^\top \in \mathbb{R}^{m+1}. \quad (14)$$

Whenever  $A$  is invertible, define

$$T(v) := A^{-1}b(\psi_*(\zeta(v | S))) \in \mathbb{R}^m. \quad (15)$$

**Auxiliary quantities.** Let  $\{(\lambda_j, \varphi_j)\}_{j=1}^n$  be an orthonormal eigendecomposition of  $L$ . For  $m \in \mathbb{N}$  and  $t > 0$ , define the truncation tail

$$\text{Tail}_t^{(m)}(u, v)^2 := \sum_{j=m+2}^n e^{-2t\lambda_j} (\varphi_j(u) - \varphi_j(v))^2. \quad (16)$$

Let  $p : V \rightarrow \mathbb{R}^m$  and  $V_1, \dots, V_{m+1} \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(V)$ . Define the degeneracy probability

$$\eta_m := \mathbb{P}(\det[p(V_1) - p(V_{m+1}) \cdots p(V_m) - p(V_{m+1})] = 0). \quad (17)$$

**Theorem 2.** Fix  $t > 0$  and  $m \in \mathbb{N}$ . Let  $G = (V, E)$  be a finite connected graph and let  $R \geq 1$ . Let  $\psi : [0, R] \rightarrow \mathbb{R}_+$  be strictly increasing. Define

$$\Delta_{\text{geom}}(G; R, t, \psi) := \sup_{\text{SPD}(u, v) \leq R} |d_t(u, v) - \psi(\text{SPD}(u, v))| \quad (18)$$

and

$$\Delta_{\text{tail}}(G; R, t, m) := \sup_{\text{SPD}(u, v) \leq R} \text{Tail}_t^{(m)}(u, v). \quad (19)$$

Then, for all  $u, v \in V$  with  $\text{SPD}(u, v) \leq R$ ,

$$|d_t^{(m)}(u, v) - \psi(\text{SPD}(u, v))| \leq \Delta_{\text{geom}}(G; R, t, \psi) + \Delta_{\text{tail}}(G; R, t, m). \quad (20)$$

For later use, define the local linkage/truncation error level

$$\delta_n^L := \Delta_{\text{geom}}(G; R, t, \psi) + \Delta_{\text{tail}}(G; R, t, m). \quad (21)$$

**Proposition 3.** Fix  $m \geq 1$  and let  $p : V \rightarrow \mathbb{R}^m$  be the truncated diffusion embedding  $p(v) := \Phi_t^{(m)}(v)$ . Draw anchors  $a_1, \dots, a_{m+1} \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(V)$  and set  $p_i := p(a_i)$ . Define

$$M := [p_1 - p_{m+1} \cdots p_m - p_{m+1}] \in \mathbb{R}^{m \times m}. \quad (22)$$

Then  $p_1, \dots, p_{m+1}$  are affinely independent in  $\mathbb{R}^m$  if and only if  $\det(M) \neq 0$ . Moreover, the trilateration matrix in Definition 1 satisfies  $A = 2M^\top$ .

Assume the degeneracy probability  $\eta_n$  in (17) satisfies  $\eta_n = o(1)$ . Then  $\mathbb{P}(\det(M) \neq 0) = 1 - \eta_n = 1 - o(1)$ .

**Remark 4** (Jittered anchors are generic a.s.). Fix  $m \geq 1$  and let  $p_1, \dots, p_{m+1} \in \mathbb{R}^m$  be arbitrary points. Let  $\varepsilon > 0$  and let  $\xi_1, \dots, \xi_{m+1} \in \mathbb{R}^m$  be random vectors whose joint law is absolutely continuous with respect to Lebesgue measure. Setting  $\tilde{p}_i := p_i + \varepsilon \xi_i$  and

$$\tilde{M} := [\tilde{p}_1 - \tilde{p}_{m+1} \cdots \tilde{p}_m - \tilde{p}_{m+1}], \quad (23)$$

we have  $\mathbb{P}(\det(\tilde{M}) \neq 0) = 1$ .

The next two theorems bound the pointwise reconstruction error  $\|\Phi_t^{(m)}(v) - T(v)\|_2$  and the induced matrix-level discrepancy.

**Theorem 5.** Let  $G \sim \mathcal{G}_{n,r}$  with fixed  $r \geq 3$ , and fix  $t > 0$  and  $m \in \mathbb{N}$ . Let  $a_1, \dots, a_{m+1}$  be anchors satisfying Proposition 3 (and, if needed, the jittered-genericity remark above), and let  $\Phi_t^{(m)}$  be the truncated spectral embedding. Suppose Assumptions 1-2 hold, and Theorem 2 holds for  $d_t^{(m)}$  with error  $\delta_n^L$ .

For a node  $v \in V$  such that  $\text{SPD}(a_i, v) \leq R$  for all  $i$ , define

$$r_i := \psi(\text{SPD}(a_i, v)), \quad r_i^* := d_t^{(m)}(a_i, v). \quad (24)$$

Then, with probability  $1 - o(1)$  over  $G$ , we have

$$\|\Phi_t^{(m)}(v) - T(v)\|_2 \leq \|A^{-1}\|_{\text{op}} \sqrt{m} (4\rho_R \delta_n^L + 2(\delta_n^L)^2), \quad (25)$$

where

$$\rho_R := \max_{0 \leq d \leq R} \psi(d) + \delta_n^L. \quad (26)$$

In particular, for  $\delta_n^L \leq 1$ ,

$$\|\Phi_t^{(m)}(v) - T(v)\|_2 \leq \|A^{-1}\|_{\text{op}} \sqrt{m} (4\rho_R + 2) \delta_n^L. \quad (27)$$

**Theorem 6.** Under the same setting and assumptions as in Theorem 5, define

$$(D_{\text{SPD}})_{v,i} = \text{SPD}(v, a_i), \quad (D_{\text{diff}}^{(m)})_{v,i} = d_t^{(m)}(v, a_i), \quad (28)$$

and let  $\psi_*(D_{\text{SPD}})$  denote the elementwise application of  $\psi$ . Then, with probability  $1 - o(1)$ ,

$$\|D_{\text{diff}}^{(m)} - \psi_*(D_{\text{SPD}})\|_F \leq \delta_n^L \sqrt{n(m+1)}. \quad (28)$$

In particular, the average per-entry discrepancy is at most  $\delta_n^L$ .

Proofs are given in Appendix A.4. Theorem 5 combines the local monotone linkage (Theorem 2) with a perturbation analysis of the linear system in Definition 1, and Theorem 6 follows by applying the linkage uniformly to all node-anchor pairs and summing the resulting errors.

### 4.3 Implications for expressivity relative to NoPE

We recall a standard consequence of distance encodings for message passing GNNs on random regular graphs.

**Theorem 7.** *Let  $G \sim \mathcal{G}_{n,r}$  be drawn from the random  $r$ -regular graph model with fixed  $r \geq 3$ , and let  $\mathcal{F}_{\text{MP}}$  denote the class of  $T$ -layer message passing GNNs without positional encodings (NoPE), whose distinguishing power is upper bounded by the 1-WL test (2019; 2018). Let  $\mathcal{F}_{\text{DE}}$  denote the class of  $T$ -layer message passing GNNs augmented with a distance encoding based on shortest-path distances to  $k = \Theta(\log n)$  anchors, as in (2009). Then there exists a family of node-level classification tasks on  $\mathcal{G}_{n,r}$  such that, with high probability over  $G$ ,*

1. *no NoPE message passing GNN in  $\mathcal{F}_{\text{MP}}$  can realize the target labels (because 1-WL cannot distinguish the relevant nodes); but*
2. *some DE-augmented GNN in  $\mathcal{F}_{\text{DE}}$  separates all label classes exactly.*

In particular, on random regular graphs the function class  $\mathcal{F}_{\text{DE}}$  is strictly more expressive than the NoPE class  $\mathcal{F}_{\text{MP}}$ .

**Corollary 1.** *Under Assumptions 1 and 2 (and the same random-regular/diffusion-geometry regime used above), consider the NoPE and DE variants of our backbone architecture: (i)  $\mathcal{F}_{\text{NoPE}}$ , the backbone without positional encodings; and (ii)  $\mathcal{F}_{\text{DE}}$ , the same backbone augmented with  $\zeta(\cdot | S)$  using  $k = \Theta(\log n)$  anchors. Then, with high probability over  $G \sim \mathcal{G}_{n,r}$  and random anchors  $S$ ,  $\mathcal{F}_{\text{DE}}$  is strictly more expressive than  $\mathcal{F}_{\text{NoPE}}$ ; in particular, there exist node classification tasks realizable by some DE-augmented instance but by no NoPE instance.*

Corollary 1 follows from Theorem 7 and the equivalence between NoPE message passing GNNs and the 1-WL test (2019; 2018). Together with Theorems 5-6, this yields the expressivity chain on random regular graphs

$$\text{NoPE} \subsetneq \text{DE-augmented} \approx \text{LapPE-augmented}, \quad (29)$$

where  $\approx$  is in the sense of the reconstruction and Frobenius-gap bounds above.

## 5 Experimental Setup

We conduct a controlled comparison of positional encodings under a unified GNP-based DDI prediction backbone. Unless stated otherwise, all implementation and hyperparameter details follow (2025b) and are provided in Appendix B.

**Datasets and protocol.** We evaluate on DrugBank and ChCh-Miner, where nodes are drugs and edges are known interactions. Molecular graphs are built from SMILES using RDKit. We adopt the inductive link prediction protocol in baselines with train/val/test splits over drug pairs.

**Baselines.** To isolate positional effects, we fix the backbone and compare three primary variants: NoPE, distance encoding (DE), and Laplacian positional encoding (LapPE). We additionally report RWSE and heat-kernel signatures (HKS) as reference positional/structural baselines under the same backbone and training protocol.

**Model configuration.** All models are implemented in PyTorch and PyTorch Geometric with standard atom/bond features. LapPE uses the first  $m$  non-trivial eigenvectors of the normalized Laplacian; DE uses shortest-path distances to  $k$  anchors with a radial transform  $\psi(\cdot)$ . RWSE uses return probabilities at steps  $\{1, 2, 4, 8, 16\}$  and HKS uses diffusion times  $\{0.1, 0.5, 1, 2, 5\}$  from a truncated Laplacian spectrum. Architecture, normalization, and optimization settings are in Appendix B.

**Evaluation.** DDI prediction is treated as binary classification; we report AUROC and F1 on the held-out test set, with the F1 threshold selected on the validation set. Primary results (NoPE/DE/LapPE) are averaged over three random seeds; RWSE/HKS are reported using the available runs.

**Ablations.** We probe the DE design by varying  $\psi(\cdot)$  on DrugBank with the number of anchors fixed, and varying the number of anchors  $k$  on ChCh-Miner with  $\psi(\cdot)$  fixed. All ablations reuse

Table 1: Validation on random  $r$ -regular graphs ( $r=6$ ,  $t=1.0$ ,  $m=8$ ; three seeds).

(a) Linkage and matrix-level discrepancy.			
$n$	$R$	$\hat{\delta}_L^n$ (mean $\pm$ std)	Frob-gap (mean $\pm$ std)
256	6	$0.1580 \pm 0.0035$	$0.0350 \pm 0.0009$
512	7	$0.1099 \pm 0.0051$	$0.0290 \pm 0.0028$
1024	7	$0.0855 \pm 0.0040$	$0.0217 \pm 0.0014$
2048	8	$0.0664 \pm 0.0024$	$0.0156 \pm 0.0002$
(b) Pointwise reconstruction and conditioning.			
$n$	$R$	$\ \Phi - T\ _2$ (median)	cond( $A$ ) (median)
256	6	0.449	46.89
512	7	0.757	112.41
1024	7	0.181	63.05
2048	8	0.0845	30.00

the main preprocessing and training settings; detailed results are in Section 6.4.

## 6 Experimental Results and Analysis

We evaluate the proposed DE-LapPE bridge from three perspectives: (i) theory-aligned validation on random  $r$ -regular graphs (Section 4); (ii) controlled spectral-approximation on real DrugBank molecular graphs; and (iii) downstream DDI prediction under a fixed GNP backbone (NoPE/DE/LapPE), together with targeted DE ablations on  $\psi(\cdot)$  and the number of anchors  $k$ .

### 6.1 Theory-aligned validation on random $r$ -regular graphs

We run a controlled study on  $G \sim \mathcal{G}_{n,r}$  with  $r = 6$  and  $n \in \{256, 512, 1024, 2048\}$  (three seeds per  $n$ ). For each graph, we form  $L = I - \frac{1}{r}A$  and compute truncated diffusion coordinates  $\Phi_t^{(m)}$  with  $t = 1.0$  and  $m = 8$ . We set the locality radius  $R = \lceil \log n \rceil$  and fit a monotone map  $\psi$  from SPD to  $d_t^{(m)}$  using isotonic regression on pairs with  $\text{SPD}(u, v) \leq R$ .

We report (i) the empirical linkage error  $\hat{\delta}_L^n := \max_{\text{SPD}(u,v) \leq R} |d_t^{(m)}(u, v) - \psi(\text{SPD}(u, v))|$ ; (ii) the normalized Frobenius gap  $\|D_{\text{diff}}^{(m)} - \psi_*(D_{\text{SPD}})\|_F / \sqrt{n(m+1)}$  on node-anchor matrices; and (iii) trilateration reconstruction with  $m+1$  anchors via  $T(v)$ , summarized by the median pointwise error  $\|\Phi_t^{(m)}(v) - T(v)\|_2$  and the median  $\text{cond}(A)$ .

Table 1 shows that both  $\hat{\delta}_L^n$  and the normalized Frobenius gap decrease as  $n$  grows, and trilateration reconstruction errors remain small in typical cases.

### 6.2 Spectral approximation of diffusion geometry on DrugBank

We analyze 80 DrugBank molecular graphs (15-200 nodes). For each graph, we compute the diffusion kernel ( $t = 1$ ) and diffusion map (top  $m = 8$  components), and approximate the kernel via an anchor-based DE Nyström scheme with  $k = 32$  anchors (farthest-point sampling) and Tikhonov regularization.

Nyström-DE is accurate: relative kernel Frobenius error  $0.024 \pm 0.021$  (median 0.020), coordinate MSE  $3.9 \times 10^{-4}$  after Procrustes alignment, and mean absolute Pearson correlation 0.988 between approximate and exact diffusion distances. Figure 3 further reports the fitted local isotonic link  $\psi^*$  (with  $R = \lceil \log n \rceil$ ) and conditioning diagnostics, showing small linkage/gap values and heavy-tailed condition numbers that motivate regularization.

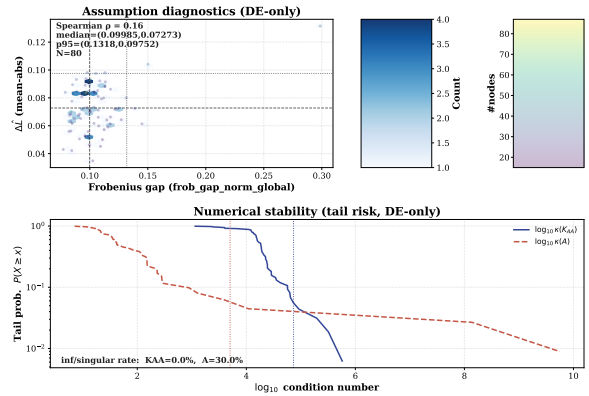


Figure 3: **DrugBank diffusion-geometry diagnostics (80 graphs)**. **Top**: normalized Frobenius gap vs. monotonicity residual  $\Delta \hat{L}$  (median/p95 and Spearman  $\rho$ ). **Bottom**: tail CCDF of  $\log_{10} \kappa(K_{AA})$  and  $\log_{10} \kappa(A)$ .

### 6.3 Effect of positional encodings on DDI prediction

Under the same GNP-DDI backbone, training schedule, and data splits, we compare NoPE, DE, LapPE, and two reference baselines (RWSE, HKS); only the positional encoding module is changed. Table 2 reports test AUROC and F1 (mean $\pm$ std over three runs).

LapPE is best overall. On DrugBank, positional information is critical: NoPE attains 0.890/0.820 (AUROC/F1), DE improves to 0.976/0.927, and LapPE further to 0.980/0.934. On ChCh-Miner, performance is already strong, but LapPE still yields consistent gains (0.946/0.879) over NoPE (0.938/0.870) and DE (0.938/0.869). RWSE is

Table 2: Test AUROC and F1 (mean  $\pm$  std over three runs) on DrugBank and ChCh-Miner under the same GNP-DDI backbone and training protocol. Only the positional encoding is varied.

Dataset	Method	Test AUROC	Test F1
DrugBank	NoPE	0.890 $\pm$ 0.002	0.820 $\pm$ 0.003
	DE	0.976 $\pm$ 0.002	0.927 $\pm$ 0.004
	LapPE	0.980 $\pm$ 0.003	0.934 $\pm$ 0.006
	RWSE	0.892 $\pm$ 0.008	0.812 $\pm$ 0.003
	HKS	0.863 $\pm$ 0.013	0.787 $\pm$ 0.007
ChCh-Miner	NoPE	0.938 $\pm$ 0.003	0.870 $\pm$ 0.002
	DE	0.938 $\pm$ 0.006	0.869 $\pm$ 0.004
	LapPE	0.946 $\pm$ 0.002	0.879 $\pm$ 0.003
	RWSE	0.944 $\pm$ 0.003	0.876 $\pm$ 0.002
	HKS	0.941 $\pm$ 0.002	0.870 $\pm$ 0.004

Table 3: Ablation study of distance encoding (DE) on DrugBank and ChCh-Miner. The best configuration for each dataset is shaded.

(a) DrugBank ( $k = 16$ , varying $\psi(\cdot)$ )					
$\psi$	Par.	Val AUC	Val F1	Test AUC	Test F1
identity	1.168	0.9430	0.8795	0.9451	0.8835
exp	1.168	0.9762	0.9279	0.9760	0.9281
log1p	1.168	0.9723	0.9196	0.9721	0.9207

(b) ChCh-Miner ( $\psi(d) = \exp(-d)$ , varying $k$ )					
$k$	Par.	Val AUC	Val F1	Test AUC	Test F1
4	0.123	0.9415	0.8715	0.9410	0.8718
8	0.123	0.9411	0.8715	0.9411	0.8727
16	0.123	0.9406	0.8709	0.9412	0.8722
32	0.124	0.9455	0.8778	0.9453	0.8784

competitive on ChCh-Miner, while HKS is weaker on DrugBank under this plug-in setting.<sup>1</sup>

## 6.4 Distance Encoding Ablation Results

On DrugBank (Table 3a), with  $k = 16$  anchors, the choice of  $\psi(\cdot)$  is decisive: identity mapping underperforms (test 0.9451/0.8835 AUROC/F1), while  $\psi = \exp$  performs best (test 0.9760/0.9281), and  $\log(1+d)$  is slightly worse. Parameter counts are identical (1.168M), so the gains come from the distance shaping rather than capacity.

On ChCh-Miner (Table 3b), fixing  $\psi(d) = \exp(-d)$ , performance is stable for  $k \in \{4, 8, 16\}$  (test  $\approx$  0.941/0.872) and improves at  $k = 32$  (test 0.9453/0.8784) with negligible parameter increase (0.123M $\rightarrow$ 0.124M).

<sup>1</sup>RWSE uses return probabilities at steps  $\{1, 2, 4, 8, 16\}$ ; HKS uses diffusion times  $\{0.1, 0.5, 1, 2, 5\}$  with truncated spectrum dimension  $k=32$  (normalized Laplacian).

## 6.5 Qualitative case study on a DrugBank molecular graph

We visualize a single DrugBank molecule (DB00006, 155 atoms) to compare the reference diffusion-map embedding from the full Gaussian-kernel eigendecomposition with the DE-based Nyström embedding using  $k = 32$  anchors (after Procrustes alignment). As shown in Fig. 4(a-b), DE Nyström closely matches the global diffusion geometry.

Fig. 4(c) reports the node-wise error  $\|\Phi_{\text{full}}(v) - \Phi_{\text{DE}}(v)\|_2$ : most nodes are at  $10^{-2}$  scale or below, with mean  $7.8 \times 10^{-3}$  and maximum  $1.15 \times 10^{-1}$ . Together with the aggregate results in Subsection 6.2, this supports that a modest anchor set can recover leading diffusion coordinates with high fidelity in practice.

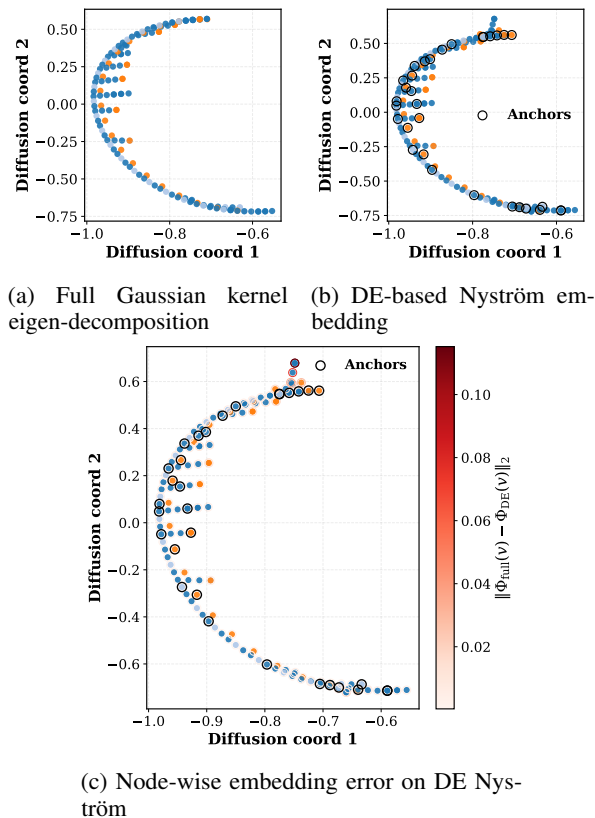


Figure 4: Qualitative comparison of diffusion-based embeddings for a single DrugBank molecular graph (DB00006).

## Limitations

We note a few limitations and practical considerations of our theory and experiments.

**Theory scope.** Our guarantees are established under a random regular graph model and a local

linkage between diffusion and shortest-path distances. They do not directly cover graphs with strong degree heterogeneity, pronounced community structure, edge weights, or directionality that commonly arise in practice.

**Approximation design choices.** The dependence of the error on truncation level, anchor placement, and the radial transform is only partially characterized. We do not analyze data-driven anchor selection, learned distance metrics, or learned (potentially non-monotone) distance transforms beyond the targeted ablations.

**Empirical generality.** We evaluate on two DDI datasets and a single multi-scale Graph Neural Process backbone; results may not fully transfer to other domains, datasets with different labeling/curation mechanisms, or architectures.

**Data and evaluation caveats.** DDI resources can be incomplete and subject to reporting/curation biases, and unobserved drug pairs used as negatives may include false negatives. As a result, performance under standard splits may not fully reflect real-world pharmacovigilance settings.

**Potential risks.** Our DDI prediction experiments are intended for research and benchmarking rather than clinical decision making. Misuse or over-reliance on predictions (false positives/negatives) could lead to inappropriate conclusions without expert review and external validation.

**Future directions.** Positional encodings are treated as fixed precomputed features; we leave joint end-to-end learning of anchors, radial maps, and spectral regularizers to future work.

**Use of AI assistants.** We used AI assistants to support code development (e.g., debugging and boilerplate). All experimental results, analyses, and claims were produced and verified by the authors.

## Acknowledgments

The authors would like to express their sincere gratitude to all the anonymous reviewers for their careful reading and insightful suggestions.

## References

Uri Alon and Eran Yahav. 2021. On the bottleneck of graph neural networks and its practical implications. In *International Conference on Learning Representations*.

Mikhail Belkin and Partha Niyogi. 2003. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396.

Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. 2017. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42.

Mengjie Chen, Ming Zhang, and Cunquan Qu. 2025. Towards interpretable drug-drug interaction prediction: A graph-based approach with molecular and network-level explanations. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pages 203–214.

Gautam Chinta, Jay Jorgenson, and Anders Karlsson. 2015. Heat kernels on regular graphs and generalized ihara zeta function formulas. *Monatshefte für Mathematik*, 178(2):171–190.

Fan Chung and Shing-Tung Yau. 1999. Coverings, heat kernels and spanning trees. *The Electronic Journal of Combinatorics*, 6(1):R12.

Ronald R Coifman and Stéphane Lafon. 2006. Diffusion maps. *Applied and computational harmonic analysis*, 21(1):5–30.

Ronald R Coifman, Stéphane Lafon, Ann B Lee, Mauro Maggioni, Boaz Nadler, Frederick Warner, and Steven W Zucker. 2005. Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. *Proceedings of the national academy of sciences*, 102(21):7426–7431.

Yifan Deng, Xinran Xu, Yang Qiu, Jingbo Xia, Wen Zhang, and Shichao Liu. 2020. A multimodal deep learning framework for predicting drug–drug interaction events. *Bioinformatics*, 36(15):4316–4322.

Vijay Prakash Dwivedi, Chaitanya K Joshi, Anh Tuan Luu, Thomas Laurent, Yoshua Bengio, and Xavier Bresson. 2023. Benchmarking graph neural networks. *Journal of Machine Learning Research*, 24(43):1–48.

Moshe Eliasof, Fabrizio Frasca, Beatrice Bevilacqua, Eran Treister, Gal Chechik, and Haggai Maron. 2023. Graph positional encoding via random feature propagation. In *International conference on machine learning*, pages 9202–9223. PMLR.

Matthias Fey and Jan Eric Lenssen. 2019. Fast graph representation learning with pytorch geometric. *Proceedings of the ICLR 2019 Workshop on Representation Learning on Graphs and Manifolds*.

Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. 2017. Neural message passing for quantum chemistry. In *International conference on machine learning*, pages 1263–1272. Pmlr.

Chao He and 1 others. 2022. Multi-type feature fusion based on graph neural network for drug–drug interaction prediction. *BMC Bioinformatics*, 23(224).

TN Kipf. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.

674	Pan Li, Yanbang Wang, Hongwei Wang, and Jure Leskovec. 2009. Distance encoding-design provably more powerful graph neural networks for structural representation learning. <i>CoRR</i> , pages 00142–00142.	728
675		729
676		
677		
678	Xuan Lin, Zhe Quan, Zhi-Jie Wang, Tengfei Ma, and Xiangxiang Zeng. 2020. Kgnn: Knowledge graph neural network for drug-drug interaction prediction. In <i>IJCAI</i> , volume 380, pages 2739–2745.	730
679		731
680		732
681		733
682	Ilya Loshchilov and Frank Hutter. 2019. Decoupled weight decay regularization. <i>International Conference on Learning Representations</i> . ICLR.	734
683		735
684		736
685	Mei Ma and Xiujuan Lei. 2023. A dual graph neural network for drug–drug interactions prediction based on molecular structure and interactions. <i>PLOS Computational Biology</i> , 19(1):e1010812.	737
686		738
687		
688		
689	Sohir Maskey, Ali Parviz, Maximilian Thiessen, Hannes Stärk, Ylli Sadikaj, and Haggai Maron. 2022. Generalized laplacian positional encoding for graph representation learning. <i>arXiv preprint arXiv:2210.15956</i> .	739
690		740
691		741
692		
693	Christopher Morris, Martin Ritzert, Matthias Fey, William L Hamilton, Jan Eric Lenssen, Gaurav Rattan, and Martin Grohe. 2019. Weisfeiler and leman go neural: Higher-order graph neural networks. In <i>Proceedings of the AAAI conference on artificial intelligence</i> , volume 33, pages 4602–4609.	742
694		743
695		744
696		745
697		
698		
699	Ding Niu, Shuai Wang, Chunyu Li, and 1 others. 2024. <a href="#">Das-ddi: A dual-view framework with drug association and similarity for drug–drug interaction prediction</a> . <i>Journal of Biomedical Informatics</i> . In press.	746
700		747
701		748
702		749
703	Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, and 1 others. 2019. Pytorch: An imperative style, high-performance deep learning library. <i>Advances in Neural Information Processing Systems</i> , 32.	750
704		751
705		752
706		753
707		
708		
709		
710	Ladislav Rampásek, Mikhail Galkin, Vijay Prakash Dwivedi, Anh Tuan Luu, Guy Wolf, and Dominique Beaini. 2022. Recipe for a general, powerful, scalable graph transformer. In <i>Advances in Neural Information Processing Systems</i> .	754
711		755
712		756
713		757
714		758
715	Jae Yong Ryu, Hyun Uk Kim, and Sang Yup Lee. 2018. Deep learning improves prediction of drug–drug and drug–food interactions. <i>Proceedings of the national academy of sciences</i> , 115(18):E4304–E4311.	759
716		760
717		761
718		762
719	Jian Sun, Maks Ovsjanikov, and Leonidas J. Guibas. 2009. A concise and provably informative multi-scale signature based on heat diffusion. In <i>Computer Graphics Forum</i> , volume 28, pages 1383–1392.	763
720		764
721		
722		
723	James Topping, Francesco Di Giovanni, Benjamin P Chamberlain, Xiaowen Dong, and Michael M Bronstein. 2022. Understanding over-squashing and bottlenecks on graphs via curvature. In <i>International Conference on Learning Representations</i> .	765
724		766
725		767
726		768
727		769
	Ulrike Von Luxburg. 2007. A tutorial on spectral clustering. <i>Statistics and computing</i> , 17(4):395–416.	770
		771
	Jiarui Wang and 1 others. 2024. Structnet-ddi: Molecular structure characterization-based representation learning for drug–drug interaction prediction. <i>Molecules</i> , 29(20):4829.	772
		773
	Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S Yu. 2020. A comprehensive survey on graph neural networks. <i>IEEE transactions on neural networks and learning systems</i> , 32(1):4–24.	774
		775
	Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2018. How powerful are graph neural networks? <i>arXiv preprint arXiv:1810.00826</i> .	776
		777
	Zimo Yan, Zheng Xie, Chang Liu, and Yuan Wang. 2025a. Resolving node identifiability in graph neural processes via laplacian spectral encodings. <i>arXiv preprint arXiv:2511.19037</i> .	778
		779
	Zimo Yan, Zheng Xie, Jie Zhang, Yiping Song, and Hao Li. 2025b. A multi-scale graph neural process with cross-drug co-attention for drug–drug interactions prediction. <i>arXiv preprint arXiv:2509.15256</i> .	780
		781
	Zimo Yan, Jie Zhang, Zheng Xie, Chang Liu, Yizhen Liu, and Yiping Song. 2025c. Metamolgen: A neural graph motif generation model for de novo molecular design. <i>arXiv preprint arXiv:2504.15587</i> .	782
		783
	Zhitao Ying, Tianjun Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and Tie-Yan Liu. 2021. Do transformers really perform badly for graph representation? In <i>Advances in Neural Information Processing Systems</i> , volume 34, pages 28877–28888.	784
		785
	Jiaxuan You, Rex Ying, Xiang Ren, William L Hamilton, and Jure Leskovec. 2019. Position-aware graph neural networks for graph-level tasks. In <i>Proceedings of the 36th International Conference on Machine Learning</i> , pages 7134–7143. PMLR. Graph-level variant of position-aware GNNs.	786
		787
	Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. 2020. Graph neural networks: A review of methods and applications. <i>AI open</i> , 1:57–81.	788
		789
	Marinka Zitnik, Monica Agrawal, and Jure Leskovec. 2018. Modeling polypharmacy side effects with graph convolutional networks. <i>Bioinformatics</i> , 34(13):i457–i466.	790
		791
		792
		793

## A Algebraic Relation Between Distance Encoding and Laplacian Spectral Coordinates

In this appendix, we give detailed proofs of the results stated in Section 4, which connect distance encoding (DE) based on shortest-path distances and Laplacian spectral coordinates via an explicit algebraic map.

Throughout, we use the notation and standing assumptions introduced in Section 3 and Section 4. In particular,  $G = (V, E)$  is a finite, simple, connected, undirected graph,  $L$  is the normalized Laplacian defined in Section 3, the truncated spectral embedding  $\Phi_t^{(m)}$  and the truncated diffusion distance  $d_t^{(m)}$  are given there, and the random-regular model, monotone linkage, anchor general position and spectral injectivity assumptions are Assumptions 1-2.

We also use the DE-to-LapPE trilateration operator  $T(\cdot)$  introduced in Definition 1, which we briefly recall for convenience.

### A.1 Reminder of the trilateration operator and distance matrices

Fix  $t > 0$  and  $m \in \mathbb{N}$ , and let  $a_1, \dots, a_{m+1} \in V$  be anchors with

$$p_i = \Phi_t^{(m)}(a_i) \in \mathbb{R}^m, \quad i = 1, \dots, m+1. \quad (30)$$

The  $m \times m$  matrix  $A$  and the map  $b : \mathbb{R}^{m+1} \rightarrow \mathbb{R}^m$  are

$$A = 2 \begin{pmatrix} (p_1 - p_{m+1})^\top \\ \vdots \\ (p_m - p_{m+1})^\top \end{pmatrix}, \quad (31)$$

$$b(r) = \begin{pmatrix} \|p_1\|_2^2 - \|p_{m+1}\|_2^2 + r_{m+1}^2 - r_1^2 \\ \vdots \\ \|p_m\|_2^2 - \|p_{m+1}\|_2^2 + r_{m+1}^2 - r_m^2 \end{pmatrix}, \quad (32)$$

for any  $r = (r_1, \dots, r_{m+1})^\top \in \mathbb{R}^{m+1}$ .

Let  $\psi : [0, R] \rightarrow \mathbb{R}_+$  be the strictly increasing function from Theorem 2. For a vector  $y = (y_1, \dots, y_{m+1})^\top \in \mathbb{R}^{m+1}$  we apply  $\psi$  elementwise and write

$$(\psi_*(y))_i = \psi(y_i), \quad i = 1, \dots, m+1. \quad (33)$$

For a node  $v \in V$  and anchor set  $S = \{a_1, \dots, a_{m+1}\}$ , recall the distance encoding

$$\zeta(v | S) = (\text{SPD}(a_1, v), \dots, \text{SPD}(a_{m+1}, v))^\top. \quad (34)$$

Whenever  $A$  is invertible, Definition 1 sets

$$T(v) = A^{-1}b(\psi_*(\zeta(v | S))) \in \mathbb{R}^m. \quad (35)$$

In addition, the node-by-anchor shortest-path and truncated diffusion distance matrices are

$$(D_{\text{SPD}})_{v,i} = \text{SPD}(v, a_i), \quad (36)$$

$$(D_{\text{diff}}^{(m)})_{v,i} = d_t^{(m)}(v, a_i) = \|\Phi_t^{(m)}(v) - \Phi_t^{(m)}(a_i)\|_2, \quad (37)$$

and we write  $\psi_*(D_{\text{SPD}})$  for the matrix obtained by applying  $\psi$  elementwise.

### A.2 Proof of Theorem 2

We begin by recalling a standard logarithmic-depth exploration window on random regular graphs. This result is proved and used as a key input in the distance-encoding analysis of Li-Wang-Wang-Leskovec (2009), and we import it here as a black-box statement in order to fix a concrete  $R = \Theta(\log n)$  regime in which local neighborhoods exhibit tree-like expansion.

**Lemma 8** (Li et al. (2009)). *Fix an integer  $r \geq 3$  and let  $G \sim \mathcal{G}_{n,r}$ . There exists a sufficiently small constant  $\epsilon > 0$  such that, with probability  $1 - o(n^{-3/2})$ , the following holds. Choose a root  $u \in V$  and define the BFS layer sets*

$$Q_k := \{v \in V : \text{SPD}(u, v) = k\}. \quad (38)$$

Let  $p_k$  denote the number of frontier half-edges incident to  $Q_k$  that remain unmatched at the beginning of the  $(k+1)$ -th BFS step in the configuration-model exposure. Then for every integer

$$k \in \left( \frac{\epsilon}{5} \cdot \frac{\log n}{\log(r-1)} + 1, \left( \frac{2}{3} - \epsilon \right) \cdot \frac{\log n}{\log(r-1)} \right), \quad (39)$$

one has

$$|Q_k| \geq (r-1-\epsilon)^{k-1} \quad \text{and} \quad p_k \geq (r-1-\epsilon)|Q_k|. \quad (40)$$

**Lemma 9.** *Let  $T_r$  be the infinite  $r$ -regular tree and let  $t > 0$ . Let  $L_{T_r}$  denote the (normalized) graph Laplacian on  $T_r$  and  $K_t := e^{-tL_{T_r}}$  its heat semigroup with heat kernel  $k_t(x, y) := (K_t)_{xy}$ . Then there exists a function  $h_t : \mathbb{N}_0 \rightarrow (0, \infty)$  such that*

$$k_t(x, y) = h_t(d(x, y)), \quad d(x, y) := \text{SPD}_{T_r}(x, y). \quad (41)$$

If one defines the diffusion distance on  $T_r$  by

$$d_t(x, y)^2 := k_{2t}(x, x) + k_{2t}(y, y) - 2k_{2t}(x, y), \quad (42)$$

then  $d_t(x, y) = \psi_{\text{tree}}(d(x, y))$  with

$$\psi_{\text{tree}}(d) := \sqrt{2(h_{2t}(0) - h_{2t}(d))}. \quad (43)$$

Moreover, for fixed  $t > 0$ ,  $h_{2t}(d)$  is strictly decreasing in  $d$ , hence  $\psi_{\text{tree}}$  is strictly increasing.

*Proof.* Since  $T_r$  is distance-transitive, for any two pairs  $(x, y)$  and  $(x', y')$  with  $d(x, y) = d(x', y')$  there exists an automorphism  $\pi$  of  $T_r$  such that  $\pi(x) = x'$  and  $\pi(y) = y'$ . The Laplacian  $L_{T_r}$  is invariant under automorphisms, hence the heat semigroup  $K_t = e^{-tL_{T_r}}$  commutes with  $\pi$ . Writing this invariance at the kernel level yields  $k_t(x, y) = k_t(\pi(x), \pi(y)) = k_t(x', y')$ , which implies the existence of a radial function  $h_t$  with  $k_t(x, y) = h_t(d(x, y))$ . Complete derivations and explicit formulas for  $h_t$  on regular trees are given in (2015); a covering-based treatment relating heat kernels on the infinite regular tree and finite regular graphs is given in (1999).

The diffusion-distance identity then follows algebraically. Since  $T_r$  is vertex-transitive,  $k_{2t}(x, x) = k_{2t}(y, y) = h_{2t}(0)$ , and radially gives  $k_{2t}(x, y) = h_{2t}(d(x, y))$ . Substituting into the definition yields

$$\begin{aligned} d_t(x, y)^2 &= h_{2t}(0) + h_{2t}(0) - 2h_{2t}(d(x, y)) \\ &= 2(h_{2t}(0) - h_{2t}(d(x, y))), \end{aligned} \quad (44)$$

so  $d_t(x, y) = \psi_{\text{tree}}(d(x, y))$  with the stated  $\psi_{\text{tree}}$ . The strict monotonicity of  $h_{2t}(d)$  in  $d$  for fixed  $t > 0$  is established in the regular-tree heat-kernel analyses cited above, which implies that  $\psi_{\text{tree}}$  is strictly increasing.  $\square$

We now prove Theorem 2. In addition to the geometric comparison hypothesis stated in the theorem, the only nontrivial step is to quantify the effect of truncating the diffusion distance to its first  $m$  nontrivial eigenmodes. We incorporate that truncation calculation directly into the proof.

*Proof of Theorem 2.* Fix  $t > 0$  and  $m \in \mathbb{N}$ . Let  $G = (V, E)$  be a finite connected graph with symmetric normalized Laplacian  $L$ . Let  $(\lambda_j, \phi_j)_{j=1}^n$  be an orthonormal eigenbasis with  $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ . For  $t > 0$ , define the full diffusion distance

$$d_t(u, v)^2 := \sum_{j=2}^n e^{-2t\lambda_j} (\phi_j(u) - \phi_j(v))^2, \quad (45)$$

and the truncated diffusion distance

$$d_t^{(m)}(u, v)^2 := \sum_{j=2}^{m+1} e^{-2t\lambda_j} (\phi_j(u) - \phi_j(v))^2. \quad (46)$$

Define the spectral tail energy

$$\text{Tail}_t^{(m)}(u, v) := \left( \sum_{j=m+2}^n e^{-2t\lambda_j} (\phi_j(u) - \phi_j(v))^2 \right)^{1/2}. \quad (47)$$

Fix any pair  $u, v \in V$  with  $\text{SPD}(u, v) \leq R$ . We first relate  $d_t^{(m)}(u, v)$  to  $d_t(u, v)$  by an exact decomposition. Splitting the defining sum of  $d_t(u, v)^2$  at index  $m + 1$  yields

$$\begin{aligned} d_t(u, v)^2 &= \sum_{j=2}^{m+1} e^{-2t\lambda_j} (\phi_j(u) - \phi_j(v))^2 \\ &\quad + \sum_{j=m+2}^n e^{-2t\lambda_j} (\phi_j(u) - \phi_j(v))^2. \end{aligned} \quad (48)$$

By definition, the first sum equals  $d_t^{(m)}(u, v)^2$  and the second sum equals  $\text{Tail}_t^{(m)}(u, v)^2$ , hence

$$d_t(u, v)^2 = d_t^{(m)}(u, v)^2 + \text{Tail}_t^{(m)}(u, v)^2. \quad (49)$$

Let  $a := d_t^{(m)}(u, v)^2 \geq 0$  and  $b := \text{Tail}_t^{(m)}(u, v)^2 \geq 0$ . Then  $d_t(u, v) = \sqrt{a+b}$  and  $d_t^{(m)}(u, v) = \sqrt{a}$ , so

$$\begin{aligned} d_t(u, v) - d_t^{(m)}(u, v) &= \sqrt{a+b} - \sqrt{a} \\ &= \frac{(a+b) - a}{\sqrt{a+b} + \sqrt{a}} = \frac{b}{\sqrt{a+b} + \sqrt{a}}. \end{aligned} \quad (50)$$

Since  $\sqrt{a+b} + \sqrt{a} \geq \sqrt{a+b} \geq \sqrt{b}$ , we obtain

$$\begin{aligned} 0 \leq d_t(u, v) - d_t^{(m)}(u, v) &= \frac{b}{\sqrt{a+b} + \sqrt{a}} \leq \frac{b}{\sqrt{b}} \\ &= \sqrt{b} = \text{Tail}_t^{(m)}(u, v), \end{aligned} \quad (51)$$

which implies

$$|d_t^{(m)}(u, v) - d_t(u, v)| \leq \text{Tail}_t^{(m)}(u, v). \quad (52)$$

We now compare  $d_t^{(m)}(u, v)$  to  $\psi(\text{SPD}(u, v))$ . By the triangle inequality,

$$\begin{aligned} |d_t^{(m)}(u, v) - \psi(\text{SPD}(u, v))| &\leq |d_t^{(m)}(u, v) - d_t(u, v)| \\ &\quad + |d_t(u, v) - \psi(\text{SPD}(u, v))|. \end{aligned} \quad (53)$$

On the high-probability event from the truncation hypothesis in the theorem statement, the first term satisfies

$$\begin{aligned} |d_t^{(m)}(u, v) - d_t(u, v)| &\leq \text{Tail}_t^{(m)}(u, v) \\ &\leq \sup_{\text{SPD}(x, y) \leq R} \text{Tail}_t^{(m)}(x, y) \leq \delta_n^{\text{trunc}}. \end{aligned} \quad (54)$$

On the high-probability event from the geometric comparison hypothesis in the theorem statement, the second term satisfies

$$|d_t(u, v) - \psi(\text{SPD}(u, v))| \leq \delta_n^{\text{geom}}. \quad (55)$$

Intersecting the two events and combining the two inequalities yields

$$|d_t^{(m)}(u, v) - \psi(\text{SPD}(u, v))| \leq \delta_n^{\text{trunc}} + \delta_n^{\text{geom}} = \delta_n^L. \quad (56)$$

Because the bounds are uniform over all pairs with  $\text{SPD}(u, v) \leq R$  on the same intersection event, the conclusion holds with high probability over  $G \sim \mathcal{G}_{n, r}$ .  $\square$

We will use Theorem 2 in the proof of the next reconstruction theorem.

### A.3 Proof of Proposition 3

*Proof of Proposition 3.* Let  $a_1, \dots, a_{m+1}$  be i.i.d. uniform on  $V$  and set  $p_i := \Phi_t^{(m)}(a_i) \in \mathbb{R}^m$ . Define  $M = [p_1 - p_{m+1} \ \dots \ p_m - p_{m+1}] \in \mathbb{R}^{m \times m}$ .

Since  $(a_1, \dots, a_{m+1})$  and  $(V_1, \dots, V_{m+1})$  have the same law, the definition (??) gives

$$\mathbb{P}(\det(M) = 0) = \eta_n. \quad (57)$$

Therefore  $\mathbb{P}(\det(M) \neq 0) = 1 - \eta_n = 1 - o(1)$ .

Assume next the jitter construction. Fix  $G$  and the anchors, hence fix  $p_1, \dots, p_{m+1}$ . Let  $\varepsilon > 0$  and let  $\xi_1, \dots, \xi_{m+1} \in \mathbb{R}^m$  be i.i.d. with a joint density, and define  $\tilde{p}_i := p_i + \varepsilon \xi_i$ . Let  $\tilde{M} = [\tilde{p}_1 - \tilde{p}_{m+1} \ \dots \ \tilde{p}_m - \tilde{p}_{m+1}]$ . Then

$$\tilde{p}_i - \tilde{p}_{m+1} = (p_i - p_{m+1}) + \varepsilon(\xi_i - \xi_{m+1}), \quad (58)$$

so

$$\tilde{M} = M + \varepsilon X, \quad X := [\xi_1 - \xi_{m+1} \ \dots \ \xi_m - \xi_{m+1}]. \quad (59)$$

Hence  $\det(\tilde{M}) = \det(M + \varepsilon X)$  is a polynomial in the entries of  $(\xi_1, \dots, \xi_{m+1})$ . This polynomial is not identically zero: indeed, choose a deterministic realization with  $\xi_{m+1} = 0$  and  $\xi_i = e_i$  for  $i =$

$1, \dots, m$ , where  $\{e_i\}$  is the standard basis of  $\mathbb{R}^m$ . Then  $X = I_m$  and

$$\det(M + \varepsilon X) = \det(M + \varepsilon I_m). \quad (60)$$

As a polynomial in  $\varepsilon$ ,  $\det(M + \varepsilon I_m)$  has leading term  $\varepsilon^m$  and therefore is not the zero polynomial, so there exist noise values making  $\det(M + \varepsilon X) \neq 0$ . Consequently the zero set

$$Z := \{(\xi_1, \dots, \xi_{m+1}) \in (\mathbb{R}^m)^{m+1} : \det(M + \varepsilon X) = 0\} \quad (61)$$

is a proper algebraic variety and has Lebesgue measure zero. Because  $(\xi_1, \dots, \xi_{m+1})$  has a joint density,

$$\mathbb{P}((\xi_1, \dots, \xi_{m+1}) \in Z \mid G, a_1, \dots, a_{m+1}) = 0, \quad (62)$$

which is equivalent to  $\mathbb{P}(\det(\tilde{M}) \neq 0 \mid G, a_1, \dots, a_{m+1}) = 1$ . This proves the affine independence under jitter.  $\square$

### A.4 Proof of Theorem 5

*Proof of Theorem 5.* Fix  $G \sim \mathcal{G}_{n, r}$ , anchors  $a_1, \dots, a_{m+1}$  and a node  $v \in V$  with  $\text{SPD}(a_i, v) \leq R$  for all  $i$ . We denote

$$z = \Phi_t^{(m)}(v) \in \mathbb{R}^m, p_i = \Phi_t^{(m)}(a_i) \in \mathbb{R}^m, \quad (63)$$

and its distance encoding with respect to  $S$  by

$$y_i = \text{SPD}(a_i, v), \quad y = \zeta(v \mid S) = (y_1, \dots, y_{m+1})^\top. \quad (64)$$

**Step 1: From shortest-path distances to approximate radii.** Let  $y_i := \text{SPD}(a_i, v)$ . By Theorem 2, there exist a strictly increasing function  $\psi$  and a sequence  $\delta_n^L = o(1)$  such that, with high probability, for all  $i$  with  $y_i \leq R$ ,

$$|d_t^{(m)}(a_i, v) - \psi(y_i)| \leq \delta_n^L. \quad (65)$$

Define

$$r_i := \psi(y_i), \quad r := (r_1, \dots, r_{m+1})^\top. \quad (66)$$

For later use, define the exact truncated radius

$$r_i^* := d_t^{(m)}(a_i, v), \quad i = 1, \dots, m+1, \quad (67)$$

and the corresponding errors

$$\zeta_{i, n} := r_i^* - r_i, \quad (68)$$

so that

$$r_i^* = r_i + \zeta_{i, n}, \quad |\zeta_{i, n}| \leq \delta_n^L. \quad (69)$$

**Step 2: Exact spectral trilateration with  $\{r_i^*\}$**

By definition of  $d_t^{(m)}$ , for each anchor  $a_i$  we have

$$d_t^{(m)}(a_i, v) = \|\Phi_t^{(m)}(a_i) - \Phi_t^{(m)}(v)\|_2 = \|p_i - z\|_2. \quad (70)$$

Squaring both sides yields

$$\|z - p_i\|_2^2 = (r_i^*)^2. \quad (71)$$

Using the identity

$$\|z - p_i\|_2^2 = \|z\|_2^2 - 2\langle z, p_i \rangle + \|p_i\|_2^2, \quad (72)$$

this becomes

$$\|z\|_2^2 - 2\langle z, p_i \rangle + \|p_i\|_2^2 = (r_i^*)^2. \quad (73)$$

These are  $m + 1$  equations in the unknown  $z \in \mathbb{R}^m$  and the scalar  $\|z\|_2^2$ .

To eliminate  $\|z\|_2^2$ , we subtract the equation corresponding to index  $m + 1$  from that corresponding to a general index  $i \in \{1, \dots, m\}$ . For each such  $i$ , we consider

$$\begin{aligned} & (\|z\|_2^2 - 2\langle z, p_i \rangle + \|p_i\|_2^2) - (\|z\|_2^2 - 2\langle z, p_{m+1} \rangle + \|p_{m+1}\|_2^2) \\ &= (r_i^*)^2 - (r_{m+1}^*)^2. \end{aligned} \quad (74)$$

On the left-hand side, we compute term by term

$$\begin{aligned} & \|z\|_2^2 - 2\langle z, p_i \rangle + \|p_i\|_2^2 - \|z\|_2^2 + 2\langle z, p_{m+1} \rangle - \|p_{m+1}\|_2^2 \\ &= -2\langle z, p_i \rangle + 2\langle z, p_{m+1} \rangle + \|p_i\|_2^2 - \|p_{m+1}\|_2^2. \end{aligned} \quad (75)$$

Factoring the inner products gives

$$-2\langle z, p_i \rangle + 2\langle z, p_{m+1} \rangle = 2\langle z, p_{m+1} - p_i \rangle = -2\langle z, p_i - p_{m+1} \rangle. \quad (76)$$

Therefore

$$-2\langle z, p_i - p_{m+1} \rangle + \|p_i\|_2^2 - \|p_{m+1}\|_2^2 = (r_i^*)^2 - (r_{m+1}^*)^2. \quad (77)$$

Multiplying both sides by  $-1$  yields

$$2\langle z, p_i - p_{m+1} \rangle = \|p_i\|_2^2 - \|p_{m+1}\|_2^2 + (r_{m+1}^*)^2 - (r_i^*)^2. \quad (78)$$

For each  $i = 1, \dots, m$ , this is a linear equation in the components of  $z$ .

We now stack these equations for  $i = 1, \dots, m$ . On the left-hand side, the  $i$ -th entry is

$$2(p_i - p_{m+1})^\top z. \quad (79)$$

Define

$$A = 2 \begin{pmatrix} (p_1 - p_{m+1})^\top \\ \vdots \\ (p_m - p_{m+1})^\top \end{pmatrix}. \quad (80)$$

Then the stacked left-hand side is exactly  $Az$ .

On the right-hand side, define

$$b^* = \begin{pmatrix} \|p_1\|_2^2 - \|p_{m+1}\|_2^2 + (r_{m+1}^*)^2 - (r_1^*)^2 \\ \vdots \\ \|p_m\|_2^2 - \|p_{m+1}\|_2^2 + (r_{m+1}^*)^2 - (r_m^*)^2 \end{pmatrix}. \quad (81)$$

The  $m$  linear equations can be written compactly as

$$Az = b^*. \quad (82)$$

By Proposition 3-??, the vectors  $p_1 - p_{m+1}, \dots, p_m - p_{m+1}$  are linearly independent in  $\mathbb{R}^m$ . Therefore the rows of  $A/2$  are linearly independent, hence  $A$  has full rank  $m$  and is invertible. There is then a unique solution

$$z = A^{-1}b^*. \quad (83)$$

Since  $z = \Phi_t^{(m)}(v)$ , this expresses the true spectral coordinate in terms of the exact truncated radii  $r_i^*$ .

**Step 3: Relating the exact and approximate right-hand sides.** In practice we construct  $T(v)$  using  $r_i$  instead of  $r_i^*$ . The corresponding right-hand side is

$$b(r) = \begin{pmatrix} \|p_1\|_2^2 - \|p_{m+1}\|_2^2 + r_{m+1}^2 - r_1^2 \\ \vdots \\ \|p_m\|_2^2 - \|p_{m+1}\|_2^2 + r_{m+1}^2 - r_m^2 \end{pmatrix}. \quad (84)$$

By Definition 1, the DE-based approximation is

$$T(v) = A^{-1}b(r). \quad (85)$$

We now compare  $b^*$  and  $b(r)$  entry by entry. For  $i = 1, \dots, m$ , we have

$$\begin{aligned} b_i^* - b(r)_i &= \left( \|p_i\|_2^2 - \|p_{m+1}\|_2^2 + (r_{m+1}^*)^2 - (r_i^*)^2 \right) \\ &\quad - \left( \|p_i\|_2^2 - \|p_{m+1}\|_2^2 + r_{m+1}^2 - r_i^2 \right) \\ &= (r_{m+1}^*)^2 - (r_i^*)^2 - r_{m+1}^2 + r_i^2. \end{aligned} \quad (86)$$

Using  $r_i^* = r_i + \zeta_{i,n}$ , we compute

$$\begin{aligned} (r_{m+1}^*)^2 - r_{m+1}^2 &= (r_{m+1} + \zeta_{m+1,n})^2 - r_{m+1}^2 \\ &= r_{m+1}^2 + 2r_{m+1}\zeta_{m+1,n} + \zeta_{m+1,n}^2 - r_{m+1}^2 \\ &= 2r_{m+1}\zeta_{m+1,n} + \zeta_{m+1,n}^2, \end{aligned} \quad (87)$$

and similarly

$$\begin{aligned} (r_i^*)^2 - r_i^2 &= (r_i + \zeta_{i,n})^2 - r_i^2 \\ &= 2r_i\zeta_{i,n} + \zeta_{i,n}^2. \end{aligned} \quad (88)$$

Substituting these expressions, we obtain

$$\begin{aligned} & (r_{m+1}^*)^2 - (r_i^*)^2 - r_{m+1}^2 + r_i^2 \\ &= ((r_{m+1}^*)^2 - r_{m+1}^2) - ((r_i^*)^2 - r_i^2) \\ &= (2r_{m+1}\zeta_{m+1,n} + \zeta_{m+1,n}^2) - (2r_i\zeta_{i,n} + \zeta_{i,n}^2). \end{aligned} \quad (89)$$

Therefore

$$b_i^* - b(r)_i = 2r_{m+1}\zeta_{m+1,n} + \zeta_{m+1,n}^2 - 2r_i\zeta_{i,n} - \zeta_{i,n}^2. \quad (90)$$

We now bound  $|b_i^* - b(r)_i|$  using the inequalities

$$|\zeta_{i,n}| \leq \delta_n^L, \quad |\zeta_{m+1,n}| \leq \delta_n^L. \quad (91)$$

Because  $m$  is fixed and the embedding dimension is finite, there exists a constant  $B_t > 0$  such that

$$\|\Phi_t^{(m)}(u)\|_2 \leq B_t \quad (92)$$

for all  $u \in V$ . Hence, for all  $i$ ,

$$r_i^* = d_t^{(m)}(a_i, v) = \|\Phi_t^{(m)}(a_i) - \Phi_t^{(m)}(v)\|_2 \leq 2B_t, \quad (93)$$

and by  $r_i = r_i^* - \zeta_{i,n}$  we have

$$|r_i| \leq |r_i^*| + |\zeta_{i,n}| \leq 2B_t + \delta_n^L. \quad (94)$$

Therefore, for all sufficiently large  $n$  (so that  $\delta_n^L \leq 1$ ), there exists a constant  $\tilde{B}_t > 0$  (depending only on  $t, m, r$ ) such that

$$|r_i| \leq \tilde{B}_t, \quad |r_{m+1}| \leq \tilde{B}_t \quad \text{for all } i = 1, \dots, m. \quad (95)$$

Using the triangle inequality, we obtain

$$\begin{aligned} |b_i^* - b(r)_i| &\leq 2|r_{m+1}| |\zeta_{m+1,n}| + \zeta_{m+1,n}^2 + 2|r_i| |\zeta_{i,n}| + \zeta_{i,n}^2 \\ &\leq 2\tilde{B}_t\delta_n^L + (\delta_n^L)^2 + 2\tilde{B}_t\delta_n^L + (\delta_n^L)^2 \\ &= 4\tilde{B}_t\delta_n^L + 2(\delta_n^L)^2. \end{aligned} \quad (96)$$

For sufficiently large  $n$ ,  $\delta_n^L \leq 1$  implies  $(\delta_n^L)^2 \leq \delta_n^L$ , hence there exists a constant  $C_1 > 0$  (depending only on  $\tilde{B}_t$ ) such that

$$|b_i^* - b(r)_i| \leq C_1 \delta_n^L, \quad i = 1, \dots, m. \quad (97)$$

Collecting the  $m$  inequalities yields

$$\|b^* - b(r)\|_2 \leq \sqrt{m} C_1 \delta_n^L. \quad (98)$$

**Step 4: Stability of the linear system and conclusion.** The true spectral coordinate and the DE-based coordinate satisfy

$$Az = b^*, \quad AT(v) = b(r). \quad (99)$$

Subtracting yields

$$A(z - T(v)) = b^* - b(r). \quad (100)$$

Since  $A$  is invertible, we have

$$z - T(v) = A^{-1}(b^* - b(r)). \quad (101)$$

Taking norms gives

$$\|z - T(v)\|_2 \leq \|A^{-1}\|_{\text{op}} \|b^* - b(r)\|_2. \quad (102)$$

Combining this with the bound on  $\|b^* - b(r)\|_2$  from Step 3 proves the claim. Recalling that  $z = \Phi_t^{(m)}(v)$  completes the proof.  $\square$

## A.5 Proof of Theorem 6

*Proof of Theorem 6.* We prove the Frobenius norm bound for the discrepancy between the truncated diffusion distance matrix  $D_{\text{diff}}^{(m)}$  and the transformed shortest-path matrix  $\psi_*(D_{\text{SPD}})$ .

Recall that

$$\begin{aligned} (D_{\text{SPD}})_{v,i} &= \text{SPD}(v, a_i), \\ (D_{\text{diff}}^{(m)})_{v,i} &= d_t^{(m)}(v, a_i), \\ (\psi_*(D_{\text{SPD}}))_{v,i} &= \psi(\text{SPD}(v, a_i)). \end{aligned} \quad (103)$$

By Theorem 2, with high probability, for all nodes  $v \in V$  and anchors  $a_i$  with  $\text{SPD}(v, a_i) \leq R$ ,

$$\left| (D_{\text{diff}}^{(m)})_{v,i} - (\psi_*(D_{\text{SPD}}))_{v,i} \right| \leq \delta_n^L. \quad (104)$$

Therefore,

$$\begin{aligned} & \left\| D_{\text{diff}}^{(m)} - \psi_*(D_{\text{SPD}}) \right\|_F^2 \\ &= \sum_{v \in V} \sum_{i=1}^{m+1} \left( (D_{\text{diff}}^{(m)})_{v,i} - (\psi_*(D_{\text{SPD}}))_{v,i} \right)^2 \\ &\leq \sum_{v \in V} \sum_{i=1}^{m+1} (\delta_n^L)^2 = n(m+1)(\delta_n^L)^2. \end{aligned} \quad (105)$$

Taking square roots yields

$$\left\| D_{\text{diff}}^{(m)} - \psi_*(D_{\text{SPD}}) \right\|_F \leq \delta_n^L \sqrt{n(m+1)}. \quad (106)$$

This proves Theorem 6.  $\square$

## B Experimental Details

### B.1 Datasets

We evaluate all methods on two widely used benchmark datasets for drug-drug interaction prediction: DrugBank and ChCh-Miner. In both cases, we follow the inductive link prediction protocol adopted in previous DDI work and in the MPNP-DDI framework (2025b).

**DrugBank.** The DrugBank interaction network is derived from the DrugBank 5.0 database and has been widely adopted in recent DDI studies (2018; 2024; 2024). The graph nodes correspond to small-molecule drugs and edges indicate documented interactions. Following standard preprocessing pipelines (2024; 2024), we obtain a graph with approximately 1,700 drugs and 190,000 labeled drug pairs, grouped into 86 interaction types in the original resource. In this work, we focus on the binary link prediction setting and treat each pair as either interacting or non-interacting. For each drug, we construct a molecular graph from its SMILES string using RDKit, where atoms are nodes with categorical and numerical attributes (such as atom type, degree, formal charge, aromaticity), and chemical bonds are edges with bond-type features.

**ChCh-Miner.** The ChCh-Miner dataset is a medium-scale DDI network curated from approved drugs and released as part of the BioSNAP collection (2018). Nodes represent drugs and edges indicate known interactions. The network contains 1,514 drugs and 48,514 documented DDI links (2024; 2024). As with DrugBank, we treat DDI prediction as a binary link prediction problem on this drug-drug graph. Molecular graphs for individual drugs are constructed from SMILES in the same way as for DrugBank.

**Data splits.** We follow the splitting strategy in MPNP-DDI (2025b). All known interactions are randomly split into training, validation, and test sets at the level of drug pairs. Negative examples are generated by uniformly sampling unobserved drug pairs, maintaining a fixed positive-to-negative ratio in each split, as is standard in DDI link prediction (2024; 2022). We adopt an inductive setting where a subset of drugs appears only in the validation or test sets, so that models must generalize to unseen molecules rather than memorizing specific pairs.

## B.2 Baselines and Model Variants

Because our focus is on understanding the role of positional encodings within a fixed architecture, we keep the backbone model identical across settings and vary only the positional encoding module. Concretely, we consider three primary variants (NoPE/DE/LapPE) for controlled comparisons, and additionally report two widely used PE baselines (RWSE and HKS) as sanity-check references under the same backbone and training protocol.

- **NoPE.** The original MPNP-DDI model (2025b) without any explicit positional encodings. Node features are given solely by atom-level descriptors derived from RDKit, and the model relies on its multi-scale message-passing scheme to infer structural information. 1163-1169
- **DE.** The backbone model augmented with distance encodings. For each molecular graph, we sample  $k$  anchor atoms and compute shortest-path distances from every node to these anchors. A radial transformation  $\psi(\cdot)$  is applied to each distance, and the resulting vectors are concatenated to the original node features. The DE module is shared across all experiments, and its design is studied in detail in the ablation experiments. 1170-1179
- **LapPE.** The backbone model augmented with Laplacian positional encodings. For each molecular graph, we compute the first  $m$  non-trivial eigenvectors of the normalized Laplacian and concatenate them to the node features, following standard practice in spectral GNNs (2023). We choose  $m$  to match the dimensionality of the DE features so that all variants have comparable parameter counts. 1180-1188
- **RWSE (reference baseline).** The backbone model augmented with random-walk structural encodings (RWSE), where each node is assigned a vector of  $K$ -step random-walk return probabilities and the resulting encoding is concatenated to node features, following common practice in graph Transformers and PE benchmarks(2022). We use steps  $\mathcal{T} = \{1, 2, 4, 8, 16\}$ . 1189-1197
- **HKS (reference baseline).** The backbone model augmented with heat-kernel signatures (HKS)(2009), constructed from a truncated eigenspace of the normalized Laplacian. Specifically, for each node we compute a diffusion-time embedding with times  $\mathcal{S} = \{0.1, 0.5, 1, 2, 5\}$  using the top- $m$  eigenpairs and concatenate it to node features. We use a truncated eigenspace dimension of  $m=32$ . 1198-1206

For fair comparison, all PE variants are used as plug-in augmentations to node features and trained with the same backbone, optimizer, and training budget. When applicable, we choose PE 1207-1210

1211 dimensionalities to be of comparable scale to the  
1212 DE/LapPE feature augmentation.

1213 All models are implemented in PyTorch (2019)  
1214 using PyTorch Geometric for efficient graph opera-  
1215 tions (2019). Unless otherwise specified, we follow  
1216 the main architectural and optimization choices of  
1217 MPNP-DDI (2025b).

### 1218 B.2.1 Feature Normalization.

1219 For molecular graphs, we construct standard atom  
1220 and bond features using RDKit, including atom  
1221 type, degree, valence, aromaticity, hybridization,  
1222 and bond type. Categorical attributes are encoded  
1223 as one-hot vectors and concatenated with numerical  
1224 descriptors. We apply feature-wise standardization  
1225 (zero mean, unit variance) across the training set for  
1226 all continuous features and reuse the same statistics  
1227 at validation and test time.

1228 For Laplacian positional encodings, raw eigen-  
1229 vectors are normalized within each molecular graph  
1230 to have zero mean and unit variance per eigenvec-  
1231 tor. To reduce sensitivity to global sign flips, we  
1232 optionally add small Gaussian noise during train-  
1233 ing. For distance encodings, shortest-path distances  
1234 are first rescaled by the median non-zero distance  
1235 within each graph and then passed through a ra-  
1236 dial map  $\psi(\cdot)$ . We consider three choices in our  
1237 experiments:  $\psi(d) = d$ ,  $\psi(d) = \exp(-d)$ , and  
1238  $\psi(d) = \log(1 + d)$ . The transformed distances are  
1239 further standardized across the training set.

### 1240 B.2.2 Architecture Design.

1241 The backbone architecture follows the multi-  
1242 scale Graph Neural Process design of MPNP-DDI  
1243 (2025b). Each drug is represented by both its  
1244 original molecular graph and its corresponding  
1245 line graph, which captures bond-level interactions.  
1246 Node and edge features (including positional en-  
1247 codings when present) are projected to a hidden  
1248 dimension of 64. The model stacks three Graph  
1249 Neural Process blocks; each block runs two itera-  
1250 tions of message passing on both the molecular  
1251 graph and the line graph, followed by aggrega-  
1252 tion and cross-scale fusion, yielding a hierarchy of  
1253 representations from local substructures to global  
1254 topology.

1255 For a pair of drugs, the block-wise representa-  
1256 tions are fed into a cross-drug co-attention mod-  
1257 ule, which computes context-aware embeddings  
1258 for each drug conditioned on its partner. These  
1259 pairwise embeddings are then passed through a  
1260 multilayer perceptron to predict the probability of

1261 an interaction. DE and LapPE variants differ only  
1262 in the additional node-level inputs; all subsequent  
1263 processing, including the co-attention and decoder,  
1264 is shared.

### 1265 B.2.3 Optimization Settings.

1266 We train all models with the AdamW optimizer  
1267 (2019) using a cosine annealing learning rate sched-  
1268 uler over 50 epochs, following (2025b). The ini-  
1269 tial learning rate and weight decay are selected  
1270 by grid search on the validation set within a stan-  
1271 dard range (for example, learning rates between  
1272  $10^{-4}$  and  $10^{-3}$  and weight decay between  $10^{-5}$   
1273 and  $10^{-3}$ ). We use binary cross-entropy loss on the  
1274 predicted interaction probabilities.

1275 Early stopping is applied based on validation  
1276 AUROC: if no improvement is observed for 10 con-  
1277 secutive epochs, training is terminated and the best  
1278 checkpoint is kept. All experiments are repeated  
1279 with three random seeds, and we report the mean  
1280 and standard deviation across runs.

### 1281 B.2.4 Batch Size and Gradient Variance.

1282 Due to the large number of drug pairs and the mem-  
1283 ory cost of the multi-scale backbone, we use a mini-  
1284 batch size of 32 drug pairs and accumulate gradi-  
1285 ents over 4 steps, resulting in an effective batch size  
1286 of 128, as in (2025b). This reduces the variance of  
1287 stochastic gradients without exceeding GPU mem-  
1288 ory constraints. We also apply gradient clipping at  
1289 a fixed maximum norm to prevent rare exploding  
1290 gradients caused by highly connected molecular  
1291 graphs.

## 1292 B.3 Evaluation Measures

1293 We treat DDI prediction as a binary classification  
1294 problem on drug pairs. For both DrugBank and  
1295 ChCh-Miner, we evaluate models using the Area  
1296 Under the Receiver Operating Characteristic Curve  
1297 (AUROC) and the F1 score, which are standard  
1298 metrics in DDI prediction (2024; 2022; 2024). AU-  
1299 ROC captures the ranking quality over all thresh-  
1300 olds, while F1 summarizes the trade-off between  
1301 precision and recall at a specific threshold.

1302 During training and model selection, we monitor  
1303 AUROC on the validation set. For F1, we select  
1304 a threshold that maximizes validation F1 and ap-  
1305 ply the same threshold to the held-out test set. All  
1306 reported numbers are computed on the test split us-  
1307 ing the checkpoint with the best validation AUROC  
1308 and are averaged over three independent runs.

#### 1309 **B.4 Details of the Ablation Study**

1310 We conduct a focused ablation study to isolate the  
1311 impact of key design choices in the distance encod-  
1312 ing module. Specifically, we examine (i) the choice  
1313 of radial transformation  $\psi(\cdot)$  applied to shortest-  
1314 path distances on DrugBank while fixing the num-  
1315 ber of anchors, and (ii) the number of anchors  $k$  on  
1316 ChCh-Miner while fixing  $\psi(\cdot)$ .

1317 On DrugBank, we fix the number of anchors  
1318 and compare three radial functions:  $\psi(d) = d$ ,  
1319  $\psi(d) = \exp(-d)$ , and  $\psi(d) = \log(1 + d)$ .  
1320 On ChCh-Miner, we fix  $\psi(d) = \exp(-d)$  and  
1321 vary  $k \in \{4, 8, 16, 32\}$  to study the trade-off  
1322 between approximation quality and model com-  
1323 plexity. All ablation runs use the same data  
1324 preprocessing, model architecture, optimizer,  
1325 batch size, and early stopping criteria as the  
1326 main experiments. Detailed numerical results  
1327 and further analysis are reported in Section 6.4.  
1328 For additional implementation details (datasets,  
1329 baselines, and hyperparameters), please re-  
1330 fer to [https://anonymous.4open.science/r/  
1331 Bridging-Distance-and-Spectral-Positional-Encodings-E48B](https://anonymous.4open.science/r/Bridging-Distance-and-Spectral-Positional-Encodings-E48B).