TANGO: Graph Neural Dynamics via Learned Energy and Tangential Flows

Extended Abstract Track Submissions

Anonymous Author(s)

Anonymous Affiliation
Anonymous Email

Abstract

We introduce TANGO, a dynamical-systems framework for graph representation learning that steers node features via a learned energy landscape. At its core is a learnable Lyapunov function whose gradient defines an energy-decreasing direction, guaranteeing stability and convergence. To preserve flexibility, we add a learned tangential message-passing component that evolves features along energy level sets. This orthogonal decomposition—gradient descent plus tangential evolution—enables effective signal propagation even in flat or ill-conditioned regions common in graph learning, mitigates oversquashing, and remains compatible with diverse GNN backbones. Empirically, TANGO achieves strong performance across node and graph classification and regression benchmarks, validating jointly learned energy functions and tangential flows.

1 Introduction

8

9

10

11 12

13

15

17

18

19

20

21

26

27

28

29

30

31

33

34

35

36

37

Graph Neural Networks (GNNs) excel on graphstructured data [12] but struggle with depth and long-range interactions due to vanishing gradients [3], over-smoothing [13, 68, 80], and over-squashing [2, 22, 39, 40, 87]. Framing GNNs as continuoustime dynamical systems (neural ODEs) [3, 15, 30, 74] enables stability analyses via diffusion [15], energy conservation [79], antisymmetric dynamics [39], and Hamiltonian flows [46]. Meanwhile, physics-informed architectures that encode conservation or dissipation improve stability and interpretability [6, 10, 36]. The common thread is reliance on an energy functional minimized or preserved by the GNN—typically simple (e.g., Dirichlet) [80]; yet many natural processes require richer energies: protein folding exhibits rugged, multi-funnel landscapes with multiple stable conformations and transition pathways [96], while complex reactions in computational chemistry demand sophisticated potential energy surfaces [82]. Energy-based models (EBMs) learn energies capturing data distributions for generative modeling [25, 42, 57, 98]. In contrast, we learn a task-driven energy whose minimization solves down-

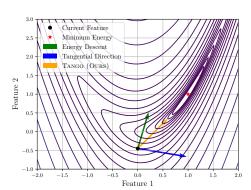


Figure 1: Illustration of TANGO dynamics in a 2D feature space. We plot level sets of a learned energy and visualize the energy-descent direction (green), the learned tangential direction (blue), and their combination (orange). The tangential component moves along level sets while the descent component reduces energy, enabling effective navigation of the learned landscape.

stream tasks (e.g., node/graph classification). This raises a central question: *How can we learn such an energy and leverage it within a GNN to guide representation dynamics?* Our answer decomposes feature evolution into two orthogonal flows: (i) a *gradient-descent* component that decreases the learned energy, and (ii) a *tangential* component that moves along its level sets, preserving energy—promoting stability, interpretability, and mitigating over-squashing.

Our Approach. We introduce TANGO, a constrained graph-dynamics framework that embeds a learnable Lyapunov energy into message passing. Updates decompose into two GNN-parameterized flows: (i) an *energy-descent* component driving convergence to task-relevant solutions, and (ii) a tangential, conservative component that preserves energy while retaining flexibility. Under mild assumptions, TANGO satisfies Lyapunov conditions, ensuring stable dynamics; the tangential flow mitigates oversquashing by enabling expressive yet controlled propagation. As illustrated in Figure 1, the descent direction (green) lowers energy, the tangential direction (blue) follows level sets, and their combination (orange) defines the full update. Across standard graph benchmarks, this structured decomposition delivers strong empirical performance—competitive with, and often surpassing, widely used baselines—while maintaining controlled, stable feature dynamics.

2 Method

43

44

45

46

47

50

52

55

56

57

58

60

62

68

71

72

78

As discussed in Section 1, our goal is to learn a task-driven energy function and use it to improve 53 downstream graph learning via TANgential- and Gradient-step Optimization of node features; we 54 therefore call our method TANGO. Section 2.1 outlines the blueprint of TANGO, Appendix B details the implementation, Section 2.2 analyzes its properties, and Appendix E discusses complexity. Mathematical background and definitions of Lyapunov functions and stability are provided in Appendix A. **Notations.** Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a graph with $n = |\mathcal{V}|$ nodes and $m = |\mathcal{E}|$ edges, and let $\mathbf{H}(t) =$ $[\mathbf{h}_1(t),\ldots,\mathbf{h}_n(t)]^{\top} \in \mathbb{R}^{n \times d}$ denote node features at continuous time t, where $\mathbf{h}_v(t) \in \mathbb{R}^d$ is the 59 state of node v. Following dynamical-systems GNNs [3, 30, 39], in discrete architectures with finitely many layers we identify time t with depth ℓ and use $\mathbf{H}(t)$ and $\mathbf{H}^{(\ell)}$ interchangeably, as appropriate.

2.1 Optimizing Features with Energy Tangential and Gradient Steps

TANGO evolves node features via a dynamical system driven by a graph energy V_G with two flows:

$$\frac{d\mathbf{H}(t)}{dt} = \underbrace{-\alpha_{\mathcal{G}}(\mathbf{H}(t)) \nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}(t))}_{\text{Energy Gradient Descent}} + \underbrace{\beta_{\mathcal{G}}(\mathbf{H}(t)) T_{V_{\mathcal{G}}}(\mathbf{H}(t))}_{\text{Tangential Direction}}, \tag{1}$$

where $\alpha_{\mathcal{G}}, \beta_{\mathcal{G}} \geq 0$ weight the flows, $\nabla_{\mathbf{H}} V_{\mathcal{G}}$ is the energy gradient, and $T_{V_{\mathcal{G}}}(\mathbf{H}(t))$ is an update direction orthogonal to it (i.e., tangential to level sets). Many orthogonal directions are possible; 65 Appendix B details how we learn T_{V_G} and implement the system. By design, the first term decreases 66 energy, while the tangential flow preserves it. 67

Tangential Flow. Setting $\beta_{\mathcal{G}} = 0$ in Equation (1) yields pure energy–gradient flow—dissipative but often slow [9, 67] and restrictive for training. While generative models tolerate hundreds-thousands of steps, this is impractical for downstream learning because it implies as many effective layers—hard to train [71] and computationally costly. To accelerate minimization, we add a tangential flow along level sets of V_G that preserves energy. As shown in Figure 1 and analyzed in Section 2.2, coupling this tangential flow with gradient descent yields a better descent direction and faster convergence.

To obtain a direction orthogonal to $\nabla_{\mathbf{H}}V_{\mathcal{G}}(\mathbf{H}(t))$, let $\mathbf{M}(\mathbf{H}(t))$ be a predicted update. Define

$$T_{V_{\mathcal{G}}}(\mathbf{H}(t)) = \mathbf{M}(\mathbf{H}(t)) - \left\langle \mathbf{M}(\mathbf{H}(t)), \widehat{\nabla}_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}(t)) \right\rangle \cdot \nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}(t)),$$
(2)

where $\widehat{\nabla}_{\mathbf{H}}V_{\mathcal{G}}(\mathbf{H}(t))$ is the normalized energy gradient. If $\nabla_{\mathbf{H}}V_{\mathcal{G}}(\mathbf{H}(t))=0$, set $T_{V_{\mathcal{G}}}(\mathbf{H}(t))=0$ $\mathbf{M}(\mathbf{H}(t))$. The projection removes the component of $\mathbf{M}(\mathbf{H}(t))$ along the descent direction, ensuring T_{V_G} is orthogonal to $\nabla_{\mathbf{H}}V_{\mathcal{G}}$. In Appendix B we specify the parameterization of TANGO as a GNN.

Theoretical Properties of TANGO

We analyze the continuous-time dynamics of TANGO in (1), focusing on (i) energy dissipation, (ii) feature evolution in flat energy landscapes, and (iii) the benefit of the tangential direction. Proofs appear in Appendix D. Throughout our analysis, we assume that (i) The input graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is 81 connected; (ii) $V_G(\mathbf{H}(t))$ is twice differentiable and bounded below. For brevity, we omit time/layer superscripts and write H for node features when clear from the context. We start by showing that 83 TANGO is dissipative if $\|\nabla_{\mathbf{H}}V_G(\mathbf{H})\|^2 > 0$, and $\alpha_G \geq 0$ (obtained by design), corresponding to the Lyapunov stability criterion from Theorem 1.

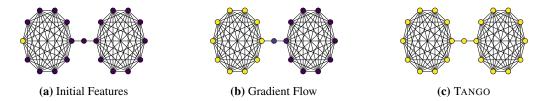


Figure 2: Comparison of propagation behaviors between gradient flow and TANGO with 50 layers. While gradient flow struggles propagating information through the bottleneck, our TANGO is effective.

Proposition 1 (Energy Dissipation). Suppose $\alpha_{\mathcal{G}} \geq 0$ and $\|\nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H})\|^2 > 0$. Then the energy $V_{\mathcal{G}}(\mathbf{H})$ is non-increasing along trajectories of Equation (1). Specifically,

$$\frac{d}{dt}V_{\mathcal{G}}(\mathbf{H}) = -\alpha_{\mathcal{G}}(\mathbf{H}) \|\nabla_{\mathbf{H}}V_{\mathcal{G}}(\mathbf{H})\|^{2} \le 0.$$
(3)

We now show that unlike gradient flows, our TANGO admits evolution of node features in flat energy landscapes, a prime challenge in optimization techniques [9, 67].

Proposition 2 (TANGO can Evolve Features in Flat Energy Landscapes). Suppose $\nabla_{\mathbf{H}}V_{\mathcal{G}}(\mathbf{H}) = 0$, and $T_{V_{\mathcal{G}}}(\mathbf{H}) \neq 0$, then the TANGO flow in Equation (1) reads: $\frac{d\mathbf{H}}{dt} = \beta_{\mathcal{G}}(\mathbf{H})T_{V_{\mathcal{G}}}(\mathbf{H})$. This implies that in contrast to gradient flows, the dynamics of TANGO can evolve even in regions where the energy landscape is flat.

91 92

93

99

100

101

104

105

106

107

114

115

119

120

121

Theoretical Benefits of Using the Tangent Direction. TANGO combines two terms (see Equations (1) and (6)): the energy gradient $\nabla_{\mathbf{H}}V_{\mathcal{G}}(\mathbf{H}^{(\ell)})$ and the tangential direction $T_{V_{\mathcal{G}}}(\mathbf{H})$. A natural question is: under what conditions does adding the tangential direction improve over plain gradient descent? We address this by first recalling a classic convergence result for gradient-based minimization.

Proposition 3 (Convergence of Gradient Descent of a Scalar Function, Nocedal and Wright [67]).

Proposition 3 (Convergence of Gradient Descent of a Scalar Function, Nocedal and Wright [67]). Let $V_{\mathcal{G}}(\cdot)$ be a scalar function and let $\mathbf{H}^{(\ell+1)} = \mathbf{H}^{(\ell)} - \alpha_{\mathcal{G}}^{(\ell)}(\mathbf{H}^{(\ell)}) \nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}^{(\ell)})$ be a gradient-descent iteration of the energy $V_{\mathcal{G}}(\cdot)$. Then, a linear convergence is obtained, with convergence rate: $r = \frac{\lambda_{\max} - \lambda_{\min}}{\lambda_{\max} + \lambda_{\min}}$, where λ_{\max} is the maximal eigenvalue, and in the case of problems that involve the graph Laplacian, λ_{\min} is the second minimal eigenvalue, i.e., the first non-zero eigenvalue of the Hessian of $V_{\mathcal{G}}(\cdot)$.

Proposition 3 shows that gradient descent struggles in ill-conditioned problems (large $\lambda_{\rm max}/\lambda_{\rm min}$), which is common in graph tasks where the Hessian inherits poor conditioning from the graph Laplacian—especially under oversquashing caused by bottlenecks [22, 38, 87]. As an alternative, we add an orthogonal flow to the gradient direction; the combined update direction is

$$\mathbf{D} = \alpha_{\mathcal{G}}(\mathbf{H}^{(\ell)}) \nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}^{(\ell)}) + \beta_{\mathcal{G}}(\mathbf{H}^{(\ell)}) T_{V_{\mathcal{G}}}(\mathbf{H}^{(\ell)}). \tag{4}$$

The following proposition demonstrates that it is possible to learn T such that \mathbf{D} becomes the Newton direction, which offers quadratic convergence [67].

Proposition 4 (TANGO can learn a Quadratic Convergence Direction). Assume for simplicity that $\beta_{\mathcal{G}} = 1$, and that the Hessian of $V_{\mathcal{G}}$ is invertible. Let $\mathbf{D} = \alpha_{\mathcal{G}}(\mathbf{H}^{(\ell)}) \nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}^{(\ell)}) + T_{V_{\mathcal{G}}}(\mathbf{H}^{(\ell)})$ with $\langle T_{V_{\mathcal{G}}}(\mathbf{H}^{(\ell)}), \widehat{\nabla}_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}^{(\ell)}) \rangle = 0$. Then, it is possible to learn a direction $T_{V_{\mathcal{G}}}(\mathbf{H}^{(\ell)})$ and a step size $\alpha_{\mathcal{G}}$ such that \mathbf{D} is the Newton direction, $\mathbf{N} = (\nabla^2 V_{\mathcal{G}})^{-1} \nabla V_{\mathcal{G}}$.

Beyond improved global convergence, Newton's method has local convergence independent of the Hessian condition number [9, 67]. Hence, if the tangential flow is learned to approximate the Newton direction, TANGO can overcome slow convergence in highly ill-conditioned landscapes—an effect well known in second-order methods and their approximations, such as conjugate gradients (CG) and L-BFGS [9, 67]. *In graph learning*, Proposition 4 is especially relevant to oversquashing [2, 22]: the graph Laplacian's smallest eigenvalue is zero (for connected graphs) and the second smallest is near zero [7, 38, 50, 87], yielding poor conditioning; under these conditions, gradient-flow methods implicitly implemented by common GNNs [23] perform poorly, limiting information propagation between nodes. By enabling feature updates that approximate second-order information, TANGO actively mitigates oversquashing. We empirically validate this in Figure 2 by comparing TANGO with Dirichlet-energy minimization commonly used by baseline GNNs [23, 80]; details are in Appendix F.

Table 1: Test performance in five benchmarks from [29]. Shown is the mean \pm _{std} of 4 runs with different random seeds. Highlighted are the top **first**, **second**, and **third** results.

Model	ZINC-12k	MNIST	CIFAR10	PATTERN	CLUSTER
	MAE↓	Accuracy [†]	Accuracy [↑]	Accuracy [†]	Accuracy [†]
GCN [52] GatedGCN [11]			$55.710_{\pm 0.381}\atop 67.312_{\pm 0.311}$		
EGT [49] GPS [75] GRIT [62]	$\begin{array}{c} 0.108_{\pm 0.009} \\ \textbf{0.070}_{\pm 0.004} \\ \textbf{0.059}_{\pm 0.002} \end{array}$	$98.051_{\pm 0.126}$	$\begin{array}{c} 68.702_{\pm 0.409} \\ \textbf{72.298}_{\pm 0.356} \\ \textbf{76.468}_{\pm 0.881} \end{array}$	$86.685_{\pm 0.059}$	
TANGO _{GatedGCN} TANGO _{GPS}			$70.894_{\pm 0.329} \\ \textbf{75.783}_{\pm 0.261}$		$78.194_{\pm 0.307} \\ \textbf{80.113}_{\pm 0.138}$

3 Experiments

125

126

127

128

130

131

132

133

134

135

137

140

141

142

143

144

145

146

148

149

150

151

155

158

159

160

161

162

We consider synthetic long-range [39], LRGB [28], Dwivedi et al. [29], and heterophilic node classification [72]. TANGO consistently improves its backbones and is competitive with other GNN architectures. Experimental details are in Appendix F, and additional results are in Appendix G.

Graph Property Prediction. We evaluate TANGO on the three graph property prediction tasks from Gravina et al. [39]—graph diameter, single-source shortest paths (SSSP), and node eccentricity on synthetic graphs—which require propagating information beyond immediate neighbors; performance thus reflects long-range interaction ability. Table 2 reports mean test $\log_{10}(MSE)$, comparing TANGO to MPNNs, DE-GNNs, and transformer-based models. Across all tasks and variants, TANGO achieves the lowest (best) error. For eccentricity, TANGO_{GPS} lowers error by > 1.2points vs. PH-DGN [46] and by > 2.0 vs. SWAN—both designed for long-radius propagation. These results validate TANGO's effectiveness at modeling long-range interactions and mitigating oversquashing. Moreover, augmenting simple backbones such as GatedGCN with TANGO consistently outperforms the baseline GatedGCN, indicating that TANGO enhances traditional MPNNs.

Table 2: Mean test set $log_{10}(MSE)(\downarrow)$ and std averaged on 4 random weight initializations on Graph Property Prediction. Lower is better. **First, second**, and **third** best results for each task are color-coded.

Model	Diameter	SSSP	Eccentricity
MPNNs			
GatedGCN [11]	$0.1348_{\pm 0.0397}$	$-3.2610_{\pm 0.0514}$	$0.6995_{\pm 0.0302}$
GCN [52]	$0.7424_{\pm0.0466}$	$0.9499_{\pm0.0001}$	$0.8468_{\pm0.0028}$
GAT [90]	$0.8221 {\scriptstyle \pm 0.0752}$	$0.6951_{\pm0.1499}$	0.7909 ± 0.0222
GraphSAGE [44]	$0.8645 {\scriptstyle \pm 0.0401}$	0.2863 ± 0.1843	0.7863 ± 0.0207
GIN [100]	$0.6131_{\pm 0.0990}$	-0.5408 ± 0.4193	$0.9504_{\pm0.0007}$
GCNII [16]	$0.5287_{\pm 0.0570}$	$-1.1329_{\pm 0.0135}$	$0.7640_{\pm 0.0355}$
DE-GNNs			
DGC [73]	$0.6028_{\pm0.0050}$	-0.1483 ± 0.0231	$0.8261_{\pm0.0032}$
GRAND [15]	$0.6715_{\pm 0.0490}$	$-0.0942_{\pm 0.3897}$	$0.6602_{\pm 0.1393}$
GraphCON [79]	$0.0964_{\pm0.0620}$	-1.3836 ± 0.0092	0.6833 ± 0.0074
A-DGN [39]	-0.5188 ± 0.1812	$-3.2417_{\pm 0.0751}$	$0.4296_{\pm0.1003}$
SWAN [40]	-0.5981 ± 0.1145		-0.0739 $_{\pm 0.2190}$
PH-DGN [46]	-0.5385 ± 0.0187	-4.2993 ± 0.0721	-0.9348 ± 0.2097
Transformers			
GPS [75]	$\text{-}0.5121 \scriptstyle{\pm 0.0426}$	-3.5990 $_{\pm 0.1949}$	$0.6077 _{\pm 0.0282}$
Ours			
TANGOGATEDGCN	$-0.6681_{\pm 0.0745}$	-5.0626 ± 0.0742	-1.7419 ± 0.0106
TANGO _{GPS}	$-0.9772_{\pm 0.0518}$	-5.5263 \pm 0.0838	-2.1455 ±0.0033

GNN Benchmarking from Dwivedi et al. [29]. To further evaluate TANGO, we use the standard GNN benchmarks from Dwivedi et al. [29], widely used for SOTA evaluation [62]. For fair comparison, we follow Dwivedi et al. [29]'s training and evaluation protocols. Table 1 reports mean±std test metrics: MAE for ZINC-12k (regression) and accuracy(%) for all others. Across all benchmarks, TANGO consistently improves its backbone and often outperforms strong baselines.

4 Conclusions

TANGO is a framework for graph neural dynamics that jointly models an energy-descent direction and a tangential flow. Casting message passing through Lyapunov theory and continuous dynamics, TANGO unifies task-driven energy-based modeling with learnable tangential flows that accelerate energy minimization. The tangential component sustains feature evolution in flat or ill-conditioned landscapes—unlike pure gradient flows—and mitigates oversquashing. Empirically, TANGO attains strong performance across 15 synthetic and real-world benchmarks, outperforming message-passing, diffusion-based, and attention-based GNNs. Future work: integrate higher-order differential operators into the tangential mechanism and develop analysis/regularization for the learned energy landscape.

References

- [1] Sami Abu-El-Haija, Bryan Perozzi, Amol Kapoor, Nazanin Alipourfard, Kristina Lerman, Hrayr Harutyunyan, Greg Ver Steeg, and Aram Galstyan. Mixhop: Higher-order graph convolutional architectures via sparsified neighborhood mixing. In *International Conference on Machine Learning*, pages 21–29. PMLR, 2019.
- [2] Uri Alon and Eran Yahav. On the bottleneck of graph neural networks and its practical implications. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=i800Ph0CVH2.
- [3] Álvaro Arroyo, Alessio Gravina, Benjamin Gutteridge, Federico Barbero, Claudio Gallicchio, Xiaowen Dong, Michael Bronstein, and Pierre Vandergheynst. On vanishing gradients, oversmoothing, and over-squashing in gnns: Bridging recurrent and graph learning. *arXiv* preprint *arXiv*:2502.10818, 2025. URL https://arxiv.org/abs/2502.10818.
 - [4] Minkyung Baek, Frank DiMaio, Ivan Anishchenko, Justas Dauparas, Sergey Ovchinnikov, Gyu Rie Lee, Jue Wang, Qian Cong, Lisa N Kinch, R Dustin Schaeffer, et al. Accurate prediction of protein structures and interactions using a three-track neural network. *Science*, 373(6557):871–876, 2021.
- [5] Ali Behrouz and Farnoosh Hashemi. Graph Mamba: Towards Learning on Graphs with State Space Models, 2024. URL https://arxiv.org/abs/2402.08678.
- [6] Ravinder Bhattoo, Sayan Ranu, and N. M. Anoop Krishnan. Learning articulated rigid body dynamics with lagrangian graph neural networks. In *Advances in Neural Information Processing Systems*, volume 35, pages 29789–29800, 2022. URL https://arxiv.org/abs/2209.11588.
 - [7] Mitchell Black, Zhengchao Wan, Amir Nayyeri, and Yusu Wang. Understanding oversquashing in gnns through the lens of effective resistance. In *Proceedings of the 40th International Conference on Machine Learning*, pages 2528–2547. PMLR, 2023.
- [8] Deyu Bo, Xiao Wang, Chuan Shi, and Huawei Shen. Beyond low-frequency information in graph convolutional networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(5):3950–3957, May 2021. doi: 10.1609/aaai.v35i5.16514. URL https://ojs.aaai.org/index.php/AAAI/article/view/16514.
- [9] Stephen P Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [10] Johannes Brandstetter, Daniel E. Worrall, and Max Welling. Message passing neural PDE solvers. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=vSix3HPYKSU.
- [11] Xavier Bresson and Thomas Laurent. Residual Gated Graph ConvNets. *arXiv preprint* arXiv:1711.07553, 2018.
 - [12] Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.
- [13] Chen Cai and Yusu Wang. A note on over-smoothing for graph neural networks. *arXiv preprint* arXiv:2006.13318, 2020.
- [14] Benjamin Chamberlain, James Rowbottom, Davide Eynard, Francesco Di Giovanni, Xiaowen
 Dong, and Michael Bronstein. Beltrami flow and neural diffusion on graphs. In *Advances in Neural Information Processing Systems*, volume 34, 2021.
- [15] Benjamin Paul Chamberlain, James Rowbottom, Maria Gorinova, Stefan Webb, Emanuele
 Rossi, and Michael M Bronstein. GRAND: Graph neural diffusion. In *International Conference* on Machine Learning (ICML), pages 1407–1418. PMLR, 2021.
- In Ing Chen, Zhewei Wei, Zengfeng Huang, Bolin Ding, and Yaliang Li. Simple and Deep Graph Convolutional Networks. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 1725–1735. PMLR, 13–18 Jul 2020.
- [17] Tian Qi Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary
 differential equations. In Advances in Neural Information Processing Systems, pages 6571–6583, 2018.

- [18] Eli Chien, Jianhao Peng, Pan Li, and Olgica Milenkovic. Adaptive universal generalized 218 pagerank graph neural network. In *International Conference on Learning Representations*, 219 2021. URL https://openreview.net/forum?id=n6jl7fLxrP.
- [19] Jeongwhan Choi, Seoyoung Hong, Noseong Park, and Sung-Bae Cho. Gread: Graph neural 221 reaction-diffusion networks. In ICML, 2023.

225

226

227

228

229

230

234

235

236

- [20] Krzysztof Choromanski, Marcin Kuczynski, Jacek Cieszkowski, Paul L. Beletsky, Konrad M. Smith, Wojciech Gajewski, Gabriel De Masson, Tomasz Z. Broniatowski, Antonina B. Gorny, Leszek M. Kaczmarek, and Stanislaw K. Andrzejewski. Performers: A new approach to scaling transformers. Proceedings of the 37th International Conference on Machine Learning (ICML), pages 2020-2031, 2020. URL https://arxiv.org/abs/2009.14743.
- [21] Krzysztof Choromanski, Han Lin, Haoxian Chen, Tianyi Zhang, Arijit Sehanobish, Valerii Likhosherstov, Jack Parker-Holder, Tamas Sarlos, Adrian Weller, and Thomas Weingarten. From block-toeplitz matrices to differential equations on graphs: towards a general theory for scalable masked transformers. In Proceedings of the 39th International Conference on Machine Learning, volume 162, pages 3962–3983. PMLR, 2022.
- [22] Francesco Di Giovanni, Lorenzo Giusti, Federico Barbero, Giulia Luise, Pietro Liò, and Michael Bronstein. On over-squashing in message passing neural networks: the impact of width, depth, and topology. In Proceedings of the 40th International Conference on Machine Learning, ICML'23. JMLR.org, 2023.
- 237 [23] Francesco Di Giovanni, James Rowbottom, Benjamin P. Chamberlain, Thomas Markovich, and Michael M. Bronstein. Graph neural networks as gradient flows. In *International Conference* 238 on Learning Representations (ICLR), 2023. URL https://arxiv.org/abs/2206.10991. 239
- [24] Lun Du, Xiaozhou Shi, Qiang Fu, Xiaojun Ma, Hengyu Liu, Shi Han, and Dongmei Zhang. 240 Gbk-gnn: Gated bi-kernel graph neural networks for modeling both homophily and heterophily. 241 In Proceedings of the ACM Web Conference 2022, WWW '22, page 1550–1558, New York, 242 NY, USA, 2022. Association for Computing Machinery. ISBN 9781450390965. doi: 10.1145/ 243 3485447.3512201. URL https://doi.org/10.1145/3485447.3512201. 244
- [25] Yilun Du and Igor Mordatch. Implicit generation and modeling with energy based models. In 245 Advances in Neural Information Processing Systems, volume 32, 2019. 246
- [26] Vijay Prakash Dwivedi and Xavier Bresson. A Generalization of Transformer Networks to 247 Graphs. AAAI Workshop on Deep Learning on Graphs: Methods and Applications, 2021. 248
- [27] Vijay Prakash Dwivedi, Anh Tuan Luu, Thomas Laurent, Yoshua Bengio, and Xavier Bresson. Graph neural networks with learnable structural and positional representations. In *International* 250 Conference on Learning Representations, 2022. URL https://openreview.net/forum? id=wTTjnvGphYj.
- [28] Vijay Prakash Dwivedi, Ladislav Rampášek, Michael Galkin, Ali Parviz, Guy Wolf, Anh Tuan 253 Luu, and Dominique Beaini. Long Range Graph Benchmark. In Advances in Neural In-254 formation Processing Systems, volume 35, pages 22326-22340. Curran Associates, Inc., 2022. 256
- [29] Vijay Prakash Dwivedi, Chaitanya K Joshi, Anh Tuan Luu, Thomas Laurent, Yoshua Bengio, 257 and Xavier Bresson. Benchmarking graph neural networks. Journal of Machine Learning 258 Research, 24(43):1–48, 2023. 259
- [30] Moshe Eliasof, Eldad Haber, and Eran Treister. PDE-GCN: Novel architectures for graph 260 neural networks motivated by partial differential equations. Advances in Neural Information 261 Processing Systems, 34:3836–3849, 2021. 262
- [31] Moshe Eliasof, Eldad Haber, and Eran Treister. Feature transportation improves graph neural 263 networks. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 38, pages 264 265 11874–11882, 2024.
- [32] Moshe Eliasof, Eldad Haber, Eran Treister, and Carola-Bibiane B Schönlieb. On the temporal 266 domain of differential equation inspired graph neural networks. In International Conference on Artificial Intelligence and Statistics, pages 1792–1800. PMLR, 2024. 268
- [33] Brandon C Fallin, Cristian F Nino, Omkar Sudhir Patil, Zachary I Bell, and Warren E Dixon. 269 Lyapunov-based graph neural networks for adaptive control of multi-agent systems. arXiv 270 preprint arXiv:2503.15360, 2025.

- Ben Finkelshtein, Xingyue Huang, Michael M. Bronstein, and Ismail Ilkan Ceylan. Cooperative Graph Neural Networks. In *Forty-first International Conference on Machine Learning*, 2024. URL https://openreview.net/forum?id=ZQcqXCuoxD.
- [35] Scott Freitas, Yuxiao Dong, Joshua Neil, and Duen Horng Chau. A large-scale database for graph representation learning. In J. Vanschoren and S. Yeung, editors, *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1, 2021.

281

282

283

284

285

286

287

291

292

293

294

295

296

297

298

299

300

301

302

303

304

305

306

307

312

313

324

- [36] Han Gao, Matthew J Zahr, and Jian-Xun Wang. Physics-informed graph neural galerkin networks: A unified framework for solving pde-governed forward and inverse problems. *Computer Methods in Applied Mechanics and Engineering*, 390:114502, 2022.
- [37] Johannes Gasteiger, Stefan Weiß enberger, and Stephan Günnemann. Diffusion Improves Graph Learning. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [38] Jhony H Giraldo, Konstantinos Skianis, Thierry Bouwmans, and Fragkiskos D Malliaros. On the trade-off between over-smoothing and over-squashing in deep graph neural networks. In *Proceedings of the 32nd ACM international conference on information and knowledge management*, pages 566–576, 2023.
- [39] Alessio Gravina, Davide Bacciu, and Claudio Gallicchio. Anti-Symmetric DGN: a stable architecture for Deep Graph Networks. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=J3Y7cgZ00S.
 - [40] Alessio Gravina, Moshe Eliasof, Claudio Gallicchio, Davide Bacciu, and Carola-Bibiane Schönlieb. On oversquashing in graph neural networks through the lens of dynamical systems. In *The 39th Annual AAAI Conference on Artificial Intelligence*, 2025.
 - [41] Albert Gu, Karan Goel, and Christopher Re. Efficiently modeling long sequences with structured state spaces. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=uYLFoz1vlAC.
 - [42] Qiushan Guo, Yifan Zhang, Yifan Wang, Yizhou Wang, and Hongsheng Li. Egc: Image generation and classification via a diffusion energy-based model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12345–12354, 2023.
 - [43] Benjamin Gutteridge, Xiaowen Dong, Michael M Bronstein, and Francesco Di Giovanni. Drew: Dynamically rewired message passing with delay. In *International Conference on Machine Learning*, pages 12252–12267. PMLR, 2023.
 - [44] William L. Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, page 1025–1035. Curran Associates Inc., 2017. ISBN 9781510860964.
 - [45] Andi Han, Dai Shi, Lequan Lin, and Junbin Gao. From continuous dynamics to graph neural networks: Neural diffusion and beyond. *arXiv preprint arXiv:2310.10121*, 2023.
- [46] Simon Heilig, Alessio Gravina, Alessandro Trenta, Claudio Gallicchio, and Davide Bacciu.

 Port-Hamiltonian Architectural Bias for Long-Range Propagation in Deep Graph Networks.

 In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=03EkqSCKuO.
 - [47] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9 (8):1735–1780, 1997. doi: 10.1162/neco.1997.9.8.1735.
- [48] Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay Pande, and Jure Leskovec. Strategies for Pre-training Graph Neural Networks. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=HJ1WWJSFDH.
- Md Shamim Hussain, Mohammed J Zaki, and Dharmashankar Subramanian. Global selfattention as a replacement for graph convolution. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 655–665, 2022.
- [50] Adarsh Jamadandi, Celia Rubio-Madrigal, and Rebekka Burkholz. Spectral graph pruning against over-squashing and over-smoothing. In *Advances in Neural Information Processing Systems*, 2024.
 - [51] Hassan K Khalil and Jessy W Grizzle. *Nonlinear systems*, volume 3. Prentice hall Upper Saddle River, NJ, 2002.

- [52] T. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks.

 Proceedings of the International Conference on Learning Representations, 2016.
- [53] Kezhi Kong, Jiuhai Chen, John Kirchenbauer, Renkun Ni, C. Bayan Bruss, and Tom Goldstein.
 GOAT: A global transformer on large-scale graphs. In Andreas Krause, Emma Brunskill,
 Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, Proceedings of the 40th International Conference on Machine Learning, volume 202 of Proceedings of Machine Learning Research, pages 17375–17390. PMLR, 23–29 Jul 2023. URL
 https://proceedings.mlr.press/v202/kong23a.html.
- Devin Kreuzer, Dominique Beaini, Will Hamilton, Vincent Létourneau, and Prudencio Tossou. Rethinking graph transformers with spectral attention. *Advances in Neural Information Processing Systems*, 34:21618–21629, 2021.
- 1337 [55] Devin Kreuzer, Dominique Beaini, Will Hamilton, Vincent Létourneau, and Prudencio Tossou.
 1338 Rethinking graph transformers with spectral attention. Advances in Neural Information
 1339 Processing Systems, 34:21618–21629, 2021.
- [56] Sven Kreuzer, Michael Reiner, and Stefan D. D. De Villiers. Sant: Structural attention networks
 for graphs. Proceedings of the 38th International Conference on Machine Learning (ICML),
 2021.
- Yann LeCun, Sumit Chopra, Raia Hadsell, Marc'Aurelio Ranzato, and Fu-Jie Huang. A tutorial on energy-based learning. *Predicting structured data*, 1(0):1–59, 2006.
- [58] Xiang Li, Renyu Zhu, Yao Cheng, Caihua Shan, Siqiang Luo, Dongsheng Li, and Weining Qian. Finding global homophily in graph neural networks when meeting heterophily. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, Proceedings of the 39th International Conference on Machine Learning, volume 162 of Proceedings of Machine Learning Research, pages 13242–13256. PMLR, 17–23 Jul 2022. URL https://proceedings.mlr.press/v162/li22ad.html.
- Daniil Likhobaba, Nikita Pavlichenko, and Dmitry Ustalov. Toloker Graph: Interaction of Crowd Annotators, February 2023. URL https://doi.org/10.5281/zenodo.7620796.

354

355

356

357

358 359

360

361

362

363

367

368

369

370

371

372 373

- [60] Meng Liu, Keqiang Yan, Bora Oztekin, and Shuiwang Ji. GraphEBM: Molecular graph generation with energy-based models. In *Energy Based Models Workshop ICLR 2021*, 2021. URL https://openreview.net/forum?id=Gc51PtL_zYw.
- [61] Sitao Luan, Chenqing Hua, Qincheng Lu, Liheng Ma, Lirong Wu, Xinyu Wang, Minkai Xu, Xiao-Wen Chang, Doina Precup, Rex Ying, Stan Z. Li, Jian Tang, Guy Wolf, and Stefanie Jegelka. The heterophilic graph learning handbook: Benchmarks, models, theoretical analysis, applications and challenges, 2024. URL https://arxiv.org/abs/2407.09618.
- [62] Liheng Ma, Chen Lin, Derek Lim, Adriana Romero-Soriano, Puneet K Dokania, Mark Coates, Philip Torr, and Ser-Nam Lim. Graph inductive biases in transformers without message passing. In *International Conference on Machine Learning*, pages 23321–23337. PMLR, 2023.
- [63] Thomas Markovich. Qdc: Quantum diffusion convolution kernels on graphs, 2023.
- [64] Sohir Maskey, Raffaele Paolino, Aras Bacho, and Gitta Kutyniok. A fractional graph laplacian approach to oversmoothing. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=kS7ED7eE74.
 - [65] Sunil Kumar Maurya, Xin Liu, and Tsuyoshi Murata. Simplifying approach to node classification in graph neural networks. *Journal of Computational Science*, 62:101695, 2022. ISSN 1877-7503. doi: https://doi.org/10.1016/j.jocs.2022.101695. URL https://www.sciencedirect.com/science/article/pii/S1877750322000990.
 - [66] Luis Müller, Mikhail Galkin, Christopher Morris, and Ladislav Rampášek. Attending to graph transformers. Transactions on Machine Learning Research, 2024. ISSN 2835-8856. URL https://openreview.net/forum?id=HhbqHBBrfZ.
 - [67] J. Nocedal and S. Wright. *Numerical Optimization*. Springer, New York, 1999.
- [68] Hoang Nt and Takanori Maehara. Revisiting graph neural networks: All we have is low-pass filters. *arXiv preprint arXiv:1905.09550*, 2019.
- Kenta Oono and Taiji Suzuki. Graph neural networks exponentially lose expressive power for node classification. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=S1ld02EFPr.

- [70] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan,
 Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. Advances in neural information processing systems, 32, 2019.
- [71] Jie Peng, Runlin Lei, and Zhewei Wei. Beyond over-smoothing: Uncovering the trainability challenges in deep graph neural networks. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pages 1878–1887, 2024.
- [72] Oleg Platonov, Denis Kuznedelev, Michael Diskin, Artem Babenko, and Liudmila Prokhorenkova. A critical look at the evaluation of GNNs under heterophily: Are we really making progress? In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=tJbbQfw-5wv.
- [73] Michael Poli, Stefano Massaroli, Junyoung Park, Atsushi Yamashita, Hajime Asama, and
 Jinkyoo Park. Graph neural ordinary differential equations. arXiv preprint arXiv:1911.07532,
 2019.
- [74] Michael Poli, Stefano Massaroli, Junyoung Park, Atsushi Yamashita, Hajime Asama, and
 Jinkyoo Park. Graph neural ordinary differential equations. arXiv preprint arXiv:1911.07532,
 2019. URL https://arxiv.org/abs/1911.07532.
- [75] Petr Rampášek, Mikhail Galkin, Vijay Prakash Dwivedi, Anh Tuan Luu, Guy Wolf, and Dominique Beaini. Recipe for a general, powerful, scalable graph transformer (graphgps). In Advances in Neural Information Processing Systems, volume 35, pages 28877–28890, 2022. URL https://arxiv.org/abs/2205.12454.
- [76] Ladislav Rampášek, Mikhail Galkin, Vijay Prakash Dwivedi, Anh Tuan Luu, Guy Wolf, and
 Dominique Beaini. Recipe for a General, Powerful, Scalable Graph Transformer. Advances in
 Neural Information Processing Systems, 35, 2022.
- Houssam Metni, Clint van Hoesel, Henrik Schopmans, Timo Sommer, et al. Graph neural networks for materials science and chemistry. *Communications Materials*, 3(1):93, 2022.
- [78] Ivan Dario Jimenez Rodriguez, Aaron Ames, and Yisong Yue. Lyanet: A lyapunov framework
 for training neural odes. In *International conference on machine learning*, pages 18687–18703.
 PMLR, 2022.
- T Konstantin Rusch, Ben Chamberlain, James Rowbottom, Siddhartha Mishra, and Michael Bronstein. Graph-coupled oscillator networks. In *International Conference on Machine Learning*, pages 18888–18909. PMLR, 2022.
- [80] T. Konstantin Rusch, Michael M. Bronstein, and Siddhartha Mishra. A Survey on Oversmoothing in Graph Neural Networks. *arXiv preprint arXiv:2303.10993*, 2023.

416

422

423

424

425

426

427

- [81] Lars Ruthotto and Eldad Haber. Deep neural networks motivated by partial differential equations. *Journal of Mathematical Imaging and Vision*, 62:352–364, 2020.
- [82] Hans M Senn and Walter Thiel. Qm/mm methods for biomolecular systems. *Angewandte Chemie International Edition*, 48(7):1198–1229, 2009.
- [83] Dai Shi, Andi Han, Lequan Lin, Yi Guo, and Junbin Gao. Exposition on over-squashing problem on gnns: Current methods, benchmarks and challenges. *arXiv preprint arXiv:2311.07073*, 2023.
 - [84] Yunsheng Shi, Zhengjie Huang, Shikun Feng, Hui Zhong, Wenjing Wang, and Yu Sun. Masked label prediction: Unified message passing model for semi-supervised classification. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 1548–1554. International Joint Conferences on Artificial Intelligence Organization, 8 2021. doi: 10.24963/ijcai.2021/214. URL https://doi.org/10.24963/ijcai.2021/214. Main Track.
 - [85] Behzad Shirzad, Amir M. Rahmani, and Marzieh Aghaei. Exphormer: Sparse attention for graphs. *Proceedings of the 40th International Conference on Machine Learning (ICML)*, 2023.
- Isolar Tönshoff, Martin Ritzert, Eran Rosenbluth, and Martin Grohe. Where did the gap go? reassessing the long-range graph benchmark. In *The Second Learning on Graphs Conference*, 2023. URL https://openreview.net/forum?id=rIUjwxc5lj.

- [87] Jake Topping, Francesco Di Giovanni, Benjamin Paul Chamberlain, Xiaowen Dong, and Michael M. Bronstein. Understanding over-squashing and bottlenecks on graphs via curvature.

 In International Conference on Learning Representations, 2022. URL https://openreview.net/forum?id=7UmjRGzp-A.
- [88] Csaba Toth, Darrick Lee, Celia Hacker, and Harald Oberhauser. Capturing graphs with hypo-elliptic diffusions. In *Advances in Neural Information Processing Systems*, 2022.
- [89] A. Vaswani et al. Attention is all you need. *Advances in Neural Information Processing*Systems, 30, 2017.
- [90] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=rJXMpikCZ.
- [91] Chloe Wang, Oleksii Tsepa, Jun Ma, and Bo Wang. Graph-mamba: Towards long-range graph sequence modeling with selective state spaces. *arXiv preprint arXiv:2402.00789*, 2024.

447

448

449

450

451

466

467

- [92] Kun Wang, Guibin Zhang, Xinnan Zhang, Junfeng Fang, Xun Wu, Guohao Li, Shirui Pan, Wei Huang, and Yuxuan Liang. The heterophilic snowflake hypothesis: Training and empowering gnns for heterophilic graphs. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '24, page 3164–3175, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400704901. doi: 10.1145/3637528. 3671791.
- 452 [93] Xiyuan Wang and Muhan Zhang. How powerful are spectral graph neural networks. In
 453 Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan
 454 Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*,
 455 volume 162 of *Proceedings of Machine Learning Research*, pages 23341–23362. PMLR,
 456 17–23 Jul 2022. URL https://proceedings.mlr.press/v162/wang22am.html.
- [94] Yifei Wang, Yisen Wang, Jiansheng Yang, and Zhouchen Lin. Dissecting the Diffusion Process
 in Linear Graph Convolutional Networks. In Advances in Neural Information Processing
 Systems, volume 34, pages 5758–5769. Curran Associates, Inc., 2021.
- Yuelin Wang, Kai Yi, Xinliang Liu, Yu Guang Wang, and Shi Jin. ACMP: Allen-cahn message passing with attractive and repulsive forces for graph neural networks. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=4fZc_79Lrqs.
- [96] Peter G Wolynes. Recent successes of the energy landscape theory of protein folding and function. *Quarterly reviews of biophysics*, 38(4):405–410, 2005.
 - [97] Louis-Pascal Xhonneux, Meng Qu, and Jian Tang. Continuous graph neural networks. In *Proceedings of the 37th International Conference on Machine Learning*, pages 10432–10441, 2020. URL https://proceedings.mlr.press/v119/xhonneux20a.html.
- [98] Jianwen Xie, Yuting Lu, Ruiqi Gao, Honglak Zhuang, and Ying Nian Wu. A theory of generative convnet. *International Conference on Machine Learning*, pages 2635–2644, 2016.
- [99] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=ryGs6iA5Km.
- [100] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=ryGs6iA5Km.
- Zhitao Ying and Jure Leskovec. Graphormer: A transformer for graphs. Proceedings of the
 43rd International ACM SIGIR Conference on Research and Development in Information
 Retrieval, 2021.
- [102] Manzil Zaheer, Guru prasad G. H., Lihong Wang, S. V. K. N. L. Wang, Yujia Li, Jakub Konečný, Shalmali Joshi, Danqi Chen, Jennifer R. R., Zhenyu Zhang, Shalini Devaraj, and Srinivas Narayanan. Bigbird: Transformers for longer sequences. *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pages 12168–12178, 2020. URL https://arxiv.org/abs/2007.14062.

- Jiong Zhu, Yujun Yan, Lingxiao Zhao, Mark Heimann, Leman Akoglu, and Danai Koutra.
 Beyond homophily in graph neural networks: Current limitations and effective designs. In
 H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, Advances in Neural
 Information Processing Systems, volume 33, pages 7793–7804. Curran Associates, Inc., 2020.
- 489 [104] Jiong Zhu, Ryan A. Rossi, Anup Rao, Tung Mai, Nedim Lipka, Nesreen K. Ahmed, and Danai 490 Koutra. Graph neural networks with heterophily. *Proceedings of the AAAI Conference on* 491 *Artificial Intelligence*, 35(12):11168–11176, May 2021. doi: 10.1609/aaai.v35i12.17332. URL 492 https://ojs.aaai.org/index.php/AAAI/article/view/17332.
- [105] Juntang Zhuang, Nicha Dvornek, Xiaoxiao Li, and James S Duncan. Ordinary differential equations on graph networks. 2020.

A Mathematical Background

495

505

515

In this section, we provide a brief overview of Lyapunov stability theory, based on the classical treatment in Khalil and Grizzle [51], which underpins the design of our TANGO. This theory originates from control systems and differential equations, offering a principled way to assess whether trajectories of a dynamical system remain bounded and converge over time.

Continuous Dynamical Systems. Let $\mathbf{h}(t) \in \mathbb{R}^d$ denote the state of a dynamical system at time $t \geq 0$, and consider a first-order ODE:

$$\frac{d\mathbf{h}(t)}{dt} = F(\mathbf{h}(t)),\tag{5}$$

where $F: \mathbb{R}^d \to \mathbb{R}^d$ is a continuous vector field. A point \mathbf{h}^* is called an *equilibrium* if $F(\mathbf{h}^*) = 0$.

Definition 1 (Lyapunov Function). Let $\mathbf{h}^* \in \mathbb{R}^d$ be an equilibrium of the system in Equation (5). A continuously differentiable function $V: \mathbb{R}^d \to \mathbb{R}$ is called a Lyapunov function around \mathbf{h}^* if:

- 1. $V(\mathbf{h}) \ge 0$ for all \mathbf{h} in a neighborhood of \mathbf{h}^* , and $V(\mathbf{h}^*) = 0$;
- 6 2. $\frac{d}{dt}V(\mathbf{h}(t)) = \nabla_{\mathbf{h}}V(\mathbf{h}(t))^{\top}F(\mathbf{h}(t)) \leq 0$ in that neighborhood.

The first condition ensures that V is lower-bounded by 0, i.e., that value of the Lyapunov function, sometimes also referred to as *energy* is non-negative, and the second that V does not increase along trajectories of the system.

We now recall a classical [51] stability criterion for the dynamical system in Equation (5), based on the definition of a Lyapunov function, which we will later use to characterize the stability of our approach in Section 2.2.

Theorem 1 (Lyapunov Stability). Let \mathbf{h}^* be an equilibrium of Equation (5) and let V be a Lyapunov function in a neighborhood \mathcal{N} of \mathbf{h}^* . If $\frac{d}{dt}V(\mathbf{h}(t)) \leq 0$ in \mathcal{N} , then \mathbf{h}^* is Lyapunov stable.

B Implementing TANGO Graph Neural Networks

As outlined in Section 2.1, TANGO is defined by a continuous dynamical system. To obtain a GNN, we discretize Equation (1) with a forward Euler step—standard in GNNs [3, 15, 19, 30, 39]—yielding the layer

$$\mathbf{H}^{(\ell+1)} = \mathbf{H}^{(\ell)} + \epsilon \left(-\alpha_{\mathcal{G}}(\mathbf{H}^{(\ell)}) \nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}^{(\ell)}) + \beta_{\mathcal{G}}(\mathbf{H}^{(\ell)}) T_{V_{\mathcal{G}}}(\mathbf{H}^{(\ell)}) \right), \tag{6}$$

for $\ell = 0, \dots, L-1$, where $\epsilon > 0$ is the Euler step size, $\nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}^{(\ell)})$ is the gradient of the energy in Equation (8), and $\alpha_{\mathcal{G}} \geq 0$, $\beta_{\mathcal{G}}$ are learned scalars that balance descent and tangential terms predicted by GNNs described below.

Energy Function. Given features $\mathbf{H}^{(\ell)}$, we implement $V_{\mathcal{G}}$ via

$$\tilde{\mathbf{H}}^{(\ell)} = \sigma \Big(\text{EnergyGNN}(\mathbf{H}^{(\ell)}; \mathcal{G}) \Big) \in \mathbb{R}^{n \times d}, \tag{7}$$

where EnergyGNN is a GNN (e.g., GatedGCN [11], GPS [75]) and σ is a pointwise nonlinearity.

Per-node energy scores (MLP) are given by $\tilde{V}_{\mathcal{G}}(\tilde{\mathbf{H}}^{(\ell)}) = \text{MLP}_{E}(\tilde{\mathbf{H}}^{(\ell)}) \in \mathbb{R}^{n \times 1}$, and the graph-level energy is defined as:

$$V_{\mathcal{G}}(\mathbf{H}^{(\ell)}) = \frac{1}{n} \sum_{v \in \mathcal{V}} \tilde{V}_{\mathcal{G}}(\tilde{\mathbf{H}}^{(\ell)})_v^2 \in \mathbb{R}_{\geq 0}.$$
 (8)

We also compute a bounded non-negative energy descent coefficient $\alpha_{\mathcal{G}}$ via global sum pooling [100] of $\tilde{\mathbf{H}}^{(\ell)}$, followed by an MLP and sigmoid:

$$\alpha_{\mathcal{G}}(\mathbf{H}^{(\ell)}) = \operatorname{Sigmoid}\left(\operatorname{MLP}_{\alpha}\left(\operatorname{SumPool}(\tilde{\mathbf{H}}^{(\ell)})\right)\right). \tag{9}$$

Tangential Update. We compute $T_{V_G}(\mathbf{H}^{(\ell)})$ with a dedicated GNN, TANGENTGNN. Given $\mathbf{H}^{(\ell)}$, it predicts a node-update direction

$$\mathbf{M}^{(\ell)} = \sigma \Big(\text{TANGENTGNN}(\mathbf{H}^{(\ell)}; \mathcal{G}) \Big). \tag{10}$$

The energy-tangential component coefficient is then obtained via the orthogonal projection in Equation (2). The tangential term is scaled by

$$\beta_{\mathcal{G}}(\mathbf{H}^{(\ell)}) = \mathrm{MLP}_{\beta}\left(\mathrm{SumPool}(\mathbf{M}^{(\ell)})\right). \tag{11}$$

C Related Work

534

535

536

537

538

541

542

548

549

550

551

552

555

556

557

562

563

564

565

566 567

569

570

578

579

580

581

582

583

584

Deep GNNs and Dynamical Systems. A growing body of work interprets GNN layers as iterative updates in a dynamical system, providing a principled framework to analyze stability, control diffusion, and inform architectural design. Poli et al. [74] introduced Graph Neural ODEs, inspired by neural ODEs [17, 81], modeling node feature evolution via continuous-depth ODEs aligned with graph structure, enabling adaptive computation and improved performance in dynamic settings. Similarly, Xhonneux et al. [97] proposed Continuous GNNs, where feature channels evolve by differential equations, mitigating over-smoothing via infinite-depth limits. Follow-up works such as GODE [105], GRAND [15], PDE-GCN_D [30], and DGC [94] view GNN layers as discrete integration steps of the heat equation to control oversmoothing [13, 68, 69]. Extensions like PDE-GCN_M [30] and GraphCON [79] add oscillatory components to preserve feature energy, while others leverage heat-kernel attention [21], anti-symmetry [39, 40], reaction-diffusion [19, 95], advection-reactiondiffusion [31] to enhance long-range or directional flow, and higher-order graph neuro ODE models [32]. A comprehensive overview is given in Han et al. [45]. Closely related, Di Giovanni et al. [23] interpret GNN layer updates as gradient flows of the Dirichlet energy, aligning message passing with energy minimization. In contrast, our TANGO learns a graph-adaptive, task-specific energy and introduces a novel descent mechanism combining energy gradients with a learnable tangential component, enabling more expressive dynamics than pure gradient flows.

Learning Energy Functions in Neural Networks. Energy-based models (EBMs) provide a flexible framework in deep learning by learning an energy function whose low-energy regions correspond to areas with high probability for the data. They have been widely used in generative tasks such as image synthesis [25, 42, 57, 98] and graph generation [60, 77]. In contrast to these typically unsupervised settings, our work focuses on learning a *task-driven* energy function tailored to predictive objectives like node or graph classification. Here, inference corresponds to descending the learned energy landscape, whose minima align with correct outputs. Relatedly, Lyapunov functions—classical tools from control theory—have been used in neural networks to ensure stable learning or inference dynamics, e.g., by enforcing stability in Neural ODEs [78] or GNN-based controllers [33]. However, such approaches typically assume a fixed or implicit energy function rather than learning one. Our method, TANGO, bridges and extends these perspectives by learning a graph-adaptive, task-specific energy and introducing a novel optimization scheme. Crucially, our TANGO incorporates a learnable tangential component that accelerates energy minimization and enhances performance in graph learning tasks.

Oversquashing in Graph Learning. Graph neural networks (GNNs) typically operate through message-passing mechanisms, aggregating information from local neighborhoods. While effective in capturing short-range dependencies, this design often leads to oversquashing, a phenomenon where signals from distant nodes are compressed into fixed-size representations, impeding the flow of longrange information [2, 22, 87]. This limitation poses a challenge in domains that demand rich global context, such as bioinformatics [4, 28] and heterophilic graphs [61, 92]. A range of strategies have been proposed to mitigate oversquashing. Graph rewiring approaches, such as SDRF [87], densify the graph to enhance connectivity prior to training. In contrast, methods like GRAND [15], BLEND [14], and DRew [43] adjust the graph structure dynamically based on node features. Transformer-based models offer another promising route by leveraging global attention to enable direct, long-range message passing. Examples include SAN [56], Graphormer [101], and GPS [76], which incorporate positional encodings, such as Laplacian eigenvectors [29] and random walk structural embeddings [27] to preserve structural identity. However, the quadratic complexity of full attention in these models raises scalability concerns, motivating interest in sparse attention mechanisms [20, 85, 102]. An alternative line of work explores *non-local dynamics* to enhance expressivity without relying solely on attention. FLODE [64] employs fractional graph operators, QDC [63] uses quantum diffusion processes, and G2TN [88] models explicit diffusion paths to propagate information more effectively. While these approaches address the oversquashing bottleneck, they often come with increased computational demands due to dense propagation operators. For a broader overview of these techniques, see Shi et al. [83]. We note that the challenge of modeling long-range dependencies also arises in other domains, such as sequential architectures [41, 47].

Optimization Techniques. The formulation of TANGO draws parallel with concepts that have been explored in the optimization literature, particularly in the design of dynamical systems that balance expressivity and convergence. While traditional gradient descent provides a robust and

interpretable mechanism for minimizing energy functions, its convergence rate can be limited in poorly conditioned settings [9, 67], which frequently arise in graph-based problems due to structural bottlenecks [2, 87]. Second-order approaches, such as Newton's method, are known to accelerate convergence by incorporating curvature information, albeit at increased computational cost. The combination of energy gradient descent and a learned tangential component in TANGO suggests a learnable departure from purely first-order schemes. Rather than explicitly computing or approximating the Hessian, our framework enables the model to learn corrective update directions that are orthogonal to the descent path. This design implicitly aligns with the motivations behind quasi-Newton techniques like conjugate gradients and LBFGS [67], which aim to improve convergence by leveraging directional information that complements the gradient. From this perspective, TANGO can be viewed as embedding optimization-inspired dynamics within graph learning frameworks. This is particularly relevant in scenarios affected by oversquashing [22], where effective feature transmission often requires departing from strictly local, gradient-driven updates. By allowing energy-preserving tangential flows, TANGO introduces flexibility reminiscent of structured optimization methods, adapted to the graph learning domain.

D Proofs of Theoretical Results

In this section, we restate the theoretical results from Section 2.2 and provide their proofs. As in the main text, we assume the following throughout: (i) the input graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is connected; (ii) the energy function $V_{\mathcal{G}}(\mathbf{H}(t))$ is twice differentiable and bounded from below. For simplicity of notation, throughout this section, we omit the time or layer scripts and use the term \mathbf{H} to denote node features when possible.

Proposition 1 (Energy Dissipation). Suppose $\alpha_{\mathcal{G}} \geq 0$ and $\|\nabla_{\mathbf{H}}V_{\mathcal{G}}(\mathbf{H})\|^2 > 0$. Then the energy $V_{\mathcal{G}}(\mathbf{H})$ is non-increasing along trajectories of Equation (1). Specifically,

$$\frac{d}{dt}V_{\mathcal{G}}(\mathbf{H}) = -\alpha_{\mathcal{G}}(\mathbf{H}) \|\nabla_{\mathbf{H}}V_{\mathcal{G}}(\mathbf{H})\|^{2} + \beta_{\mathcal{G}}(\mathbf{H})\langle T_{V_{\mathcal{G}}}(\mathbf{H}), \nabla_{\mathbf{H}}V_{\mathcal{G}}(\mathbf{H})\rangle$$

$$= -\alpha_{\mathcal{G}}(\mathbf{H}) \|\nabla_{\mathbf{H}}V_{\mathcal{G}}(\mathbf{H})\|^{2} \leq 0.$$

611 *Proof.* By the chain rule,

$$\frac{d}{dt}V_{\mathcal{G}}(\mathbf{H}) = \left\langle \nabla_{\mathbf{H}}V_{\mathcal{G}}(\mathbf{H}), \frac{d\mathbf{H}}{dt} \right\rangle.$$

Substituting the dynamics of Equation (1):

$$\frac{d}{dt}V_{\mathcal{G}}(\mathbf{H}) = \langle \nabla_{\mathbf{H}}V_{\mathcal{G}}(\mathbf{H}), -\alpha_{\mathcal{G}}(\mathbf{H}) \nabla_{\mathbf{H}}V_{\mathcal{G}}(\mathbf{H}) + \beta_{\mathcal{G}}(\mathbf{H}) T_{V_{\mathcal{G}}}(\mathbf{H}) \rangle
= -\alpha_{\mathcal{G}}(\mathbf{H}) \|\nabla_{\mathbf{H}}V_{\mathcal{G}}(\mathbf{H})\|^{2} + \beta_{\mathcal{G}}(\mathbf{H}) \langle T_{V_{\mathcal{G}}}(\mathbf{H}), \nabla_{\mathbf{H}}V_{\mathcal{G}}(\mathbf{H}) \rangle.$$

As discussed in Section 2, we have by design, that

$$\langle T_{V_{\mathcal{G}}}(\mathbf{H}), \nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}) \rangle = 0.$$

Therefore,

588

589

590

591

593

595

596

597

598

599

603

607

608

$$\frac{d}{dt}V_{\mathcal{G}}(\mathbf{H}) = -\alpha_{\mathcal{G}}(\mathbf{H}) \|\nabla_{\mathbf{H}}V_{\mathcal{G}}(\mathbf{H})\|^{2}.$$

Because $\alpha_{\mathcal{G}}(\mathbf{H}) \geq 0$ by design, the energy is non-increasing, and assuming $\alpha_{\mathcal{G}}(\mathbf{H}) > 0$, the system is dissipative, i.e., its energy is decreasing.

Proposition 2 (TANGO can Evolve Features in Flat Energy Landscapes). Suppose $\nabla_{\mathbf{H}}V_{\mathcal{G}}(\mathbf{H}) = 0$, and $T_{V_{\mathcal{G}}}(\mathbf{H}) \neq 0$, then the TANGO flow in Equation (1) reads:

$$\frac{d\mathbf{H}}{dt} = \beta_{\mathcal{G}}(\mathbf{H}) T_{V_{\mathcal{G}}}(\mathbf{H}).$$

This implies that in contrast to gradient flows, the dynamics of TANGO can evolve even in regions where the energy landscape is flat.

Proof. Because $\nabla_{\mathbf{H}}V_{\mathcal{G}}(\mathbf{H})=0$, the first term in Equation (1) vanishes, and the TANGO dynamical system reads:

$$\frac{d\mathbf{H}}{dt} = \beta_{\mathcal{G}}(\mathbf{H}) T_{V_{\mathcal{G}}}(\mathbf{H}),$$

- Assuming that $T_{V_{\mathcal{G}}}(\mathbf{H}) \neq 0$, TANGO can continue evolving node features also in cases where $\nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}) = 0$, i.e., where the energy landscape is flat.
- Proposition 3 (Convergence of Gradient Descent of a Scalar Function, Nocedal and Wright [67]). Let $V_{\mathcal{G}}(\cdot)$ be a scalar function and let $\mathbf{H}^{(\ell+1)} = \mathbf{H}^{(\ell)} \alpha_{\mathcal{G}}^{(\ell)}(\mathbf{H}^{(\ell)}) \nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}^{(l)})$ be a gradient-descent
- iteration of the energy $V_{\mathcal{G}}(\cdot)$. Then, a linear convergence is obtained, with convergence rate:

$$r = \frac{\lambda_{\text{max}} - \lambda_{\text{min}}}{\lambda_{\text{max}} + \lambda_{\text{min}}},$$

- where λ_{max} is the maximal eigenvalue, and in the case of problems that involve the graph Laplacian,
- λ_{\min} is the second minimal eigenvalue, i.e., the first non-zero eigenvalue of the Hessian of $V_{\mathcal{G}}(\cdot)$.
- Proposition 4 (TANGO can learn a Quadratic Convergence Direction). Assume for simplicity that
- $\beta_{\mathcal{G}} = 1, \text{ and that the Hessian of } V_{\mathcal{G}} \text{ is invertible. Let } \mathbf{D} = \alpha_{\mathcal{G}}(\mathbf{H}^{(\ell)}) \nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}^{(\ell)}) + T_{V_{\mathcal{G}}}(\mathbf{H}^{(\ell)}) \text{ with}$
- $\left\langle T_{V_{\mathcal{G}}}(\mathbf{H}^{(\ell)}), \widehat{\nabla}_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}^{(\ell)}) \right\rangle = 0. \text{ Then, it is possible to learn a direction } T_{V_{\mathcal{G}}}(\mathbf{H}^{(\ell)}) \text{ and a step size } T_{V_{\mathcal{G}}}(\mathbf{H}^{(\ell)})$
- 631 $\alpha_{\mathcal{G}}$ such that **D** is the Newton direction, $\mathbf{N} = (\nabla^2 V_{\mathcal{G}})^{-1} \nabla V_{\mathcal{G}}$.
- Proof. We aim to construct a direction $\mathbf{D} = \alpha_{\mathcal{G}}(\mathbf{H}) \nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}) + T_{V_{\mathcal{G}}}(\mathbf{H})$ that matches the Newton direction:

$$\mathbf{N} = \left(\nabla_{\mathbf{H}}^2 V_{\mathcal{G}}(\mathbf{H})\right)^{-1} \nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}).$$

Recall that by design, we have that $T_{V_G}(\mathbf{H})$ is orthogonal to the energy gradient, i.e., $\langle T_{V_G}(\mathbf{H}), \nabla_{\mathbf{H}} V_G(\mathbf{H}) \rangle = 0$. Then, we can express a Newton direction by the decomposition:

$$\mathbf{N} = \alpha_{\mathcal{G}}(\mathbf{H}) \, \nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}) + T_{V_{\mathcal{G}}}(\mathbf{H}).$$

636 Solving for the orthogonal component yields:

$$T_{V_{\mathcal{G}}}(\mathbf{H}) = \mathbf{N} - \alpha_{\mathcal{G}}(\mathbf{H}) \nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}).$$

To enforce orthogonality, we require:

$$\langle \mathbf{N} - \alpha_{\mathcal{G}}(\mathbf{H}) \nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}), \nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}) \rangle = 0.$$

Expanding and simplifying, we find:

$$\langle \mathbf{N}, \nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}) \rangle - \alpha_{\mathcal{G}}(\mathbf{H}) \| \nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}) \|^2 = 0,$$

and the optimal step size is given by:

641

642

$$\alpha_{\mathcal{G}}(\mathbf{H}) = \frac{\langle \mathbf{N}, \nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}) \rangle}{\left\| \nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}) \right\|^2},$$

showing that it is possible to learn a Newton direction, i.e., a quadratic energy convergence direction.

E Complexity and Runtimes

Complexity. Each step of TANGO requires computing the gradient of the learned energy function 643 $V_G(\mathbf{H}^{(\ell)})$, that is defined in Equation (8). This involves two main operations: (i) forward and 644 backward passes through the energy network ENERGYGNN, which contains L_{energy} message-passing layers and an MLP; and (ii) automatic differentiation to compute $\nabla_{\mathbf{H}}V_{\mathcal{G}}(\mathbf{H}^{(\ell)})$ with respect to the input node features. In parallel, the tangential flow direction $T_{V_G}(\mathbf{H}^{(\ell)})$ is obtained by projecting the 647 vector field $\mathbf{M}^{(\ell)}$ computed by a separate TANGENTGNN with L_{tangent} layers onto the orthogonal 648 complement of the normalized energy gradient, as shown in Equation (2). This projection is of 649 computational cost of O(nd) per step, where $n = |\mathcal{V}|$ and d is the feature dimensionality. In addition, 650 scalar coefficients α_G and β_G are computed from pooled node features using MLPs (Equations (9) and (11)). Assuming both ENERGYGNN and TANGENTGNN are message-passing architectures

with linear complexity in the number of nodes and edges, and setting $L_{\text{energy}} = L_{\text{tangent}}$, the total complexity per layer becomes $O(L_{\text{gnn}} \cdot (n+m) \cdot d)$, where L_{gnn} is the number of GNN layers used in each subnetwork and $m = |\mathcal{E}|$ is the number of edges. Unrolling the dynamics over L steps, the overall computational complexity of TANGO is:

$$O(L \cdot L_{gnn} \cdot (|\mathcal{V}| + |\mathcal{E}|) \cdot d)$$
.

Runtimes. We measure the runtimes of our TANGO using two backbones, GatedGCN and GPS, and compare it with the baseline backbone runtimes. In addition, we consider other methods like FAGCN [8] and CO-GNN [34] for a broad comparison of the runtimes of TANGO. For reference, we also refer to Table 7, where we compare the obtained downstream performance, which shows in many cases significant improvement using our TANGO variants compared with other considered methods. We measure the runtimes on the Questions dataset, using the same major hyperparameters for all methods (256 channels, 8 layers) to ensure fairness. The measurements were conducted on an NVIDIA RTX6000 Ada GPU with 48GB of memory.

Table 3: Training runtimes (milliseconds per epoch) on the Questions dataset using an 8-layer network with 256 channels on an NVIDIA RTX6000 Ada GPU.

Method GCN	CO-GNN	FAGCN	GatedGCN	GAT	GPS(GatedGCN)	TANGO _{GatedGCN}	TANGO _{GPS}
Runtime (ms/epoch) 69.77	210.32	103.94	129.92	112.40	429.08	184.98	694.27

F Experimental Details

665

667

668

669

672

674

676

677

678

680

681

682

685

686

687

688

689

In this section, we provide additional experimental details.

Computational Resources. Our experiments are run on NVIDIA RTX6000 Ada with 48GB of memory. Our code is implemented in PyTorch [70], and will be publicly released upon acceptance.

Baselines. We consider different classical and state-of-the-art GNN baselines. Specifically:

- Classical MPNNs, i.e., GCN [52], GraphSAGE [44], GAT [90], GatedGCN [11], GIN [99], GINE [48], GCNII [16], and CoGNN [34];
- Heterophily-specific models, i.e., H2GCN [103], CPGNN [104], FAGCN [8], GPR-GNN [18], FSGNN [65], GloGNN [58], GBK-GNN [24], and JacobiConv [93];
- DE-DGNs, i.e., DGC [94], GRAND [15], GraphCON [79], A-DGN [39], and SWAN [40];
- Graph Transformers, i.e., Transformer [26, 89], GT [84], SAN [55], GPS [76], GOAT [53], and Exphormer [85];
- Higher-Order DGNs, i.e., DIGL [37], MixHop [1], and DRew [43].
- SSM-based GNN, i.e., Graph-Mamba [91], GMN [5], and GPS+Mamba [5]

F.1 Synthetic Example from Figure 2

In the synthetic example in Figure 2, we demonstrate the effectiveness of TANGO in overcoming the oversquashing issue in GNNs. To do that, we consider a Barbell graph, where all node features are set to 0, besides the left-most node in the graph, which is set to 1, as shown in Figure 2(a). The goal is to allow the information to propagate through all nodes effectively. We do this by considering a gradient flow process of the Dirichlet energy using 50 layers (steps), as shown in Figure 2(b), where it is noticeable that the information is now flowing to the right part in the graph, because of the bottleneck between the two cliques. However, as we show in Figure 2(c), by considering our TANGO, which utilizes both an energy flow as well as a tangential flow, it is possible to effectively propagate the information through all the nodes in the graphs.

F.2 Graph Property Prediction

Dataset. We construct our benchmark following the protocol introduced by Gravina et al. [39]. Graph instances are synthetically generated from a variety of canonical topologies, including Erdős–Rényi, Barabasi-Albert, caveman, tree, and grid models. Each graph consists of 25 to 35 nodes, with

node features initialized as random identifiers sampled uniformly from the interval [0, 1). The prediction targets encompass several structural tasks: computing the shortest paths from a source node, estimating node eccentricity, and determining graph diameter. The complete dataset contains 7,040 graphs, split into 5,120 for training, 640 for validation, and 1,280 for testing. These tasks inherently demand capturing long-range dependencies, as they involve global graph computations such as shortest path inference. As highlighted in Gravina et al. [39], traditional algorithms like Bellman-Ford or Dijkstra's method require multiple rounds of message propagation, which motivates the need for expressive graph models. The benchmark graph families, such as caveman, tree, line, star, caterpillar, and lobster, frequently include structural bottlenecks that are known to induce oversquashing effects [87], posing additional challenges for message-passing-based GNNs.

Experimental Setup. We adopt the same evaluation framework as Gravina et al. [39], including datasets, training routines, and hyperparameter spaces. Model training is conducted using the Adam optimizer for up to 1500 epochs, with early stopping triggered after 100 consecutive epochs of no improvement on the validation Mean Squared Error (MSE). Hyperparameters are selected via grid search, and performance is averaged over 4 independent runs with different random seeds for weight initialization. A summary of the hyperparameter grid used in our experiments is provided in Table 5.

F.3 Graph Benchmarks from Dwivedi et al. [29]

694

695

696

697

698

700

701

703

708

709

712

713

714

715

716

717

720

726

730

736

740

741

Dataset. To comprehensively assess the capabilities of TANGO, we evaluate its performance on a diverse set of graph learning benchmarks curated by Dwivedi et al. [29]. The benchmark suite includes: ZINC-12k, a molecular regression dataset containing chemical compounds, where the goal is to predict the constrained solubility of each molecule. Graphs represent molecular structures, with atoms as nodes and chemical bonds as edges. Node and edge features encode atom types and bond types, respectively. MNIST and CIFAR-10 superpixels are graph-structured versions of standard image classification datasets, where images are converted into sparse graphs of superpixels. Each superpixel forms a node, and edges are based on spatial adjacency. The tasks involve classifying digits (MNIST) and natural objects (CIFAR-10) based on graph-structured representations. *CLUSTER* and PATTERN are synthetic datasets designed to assess the relational inductive biases of graph neural networks. Both datasets are generated from a set of stochastic block models (SBMs). In *CLUSTER*, the task is to group nodes by community, while PATTERN involves identifying specific structural patterns within each graph. These datasets span a variety of domains: chemical, image, and synthetic graphs, and are commonly used to benchmark architectural innovations in GNNs [62]. We follow the official training, validation, and test splits provided by Dwivedi et al. [29], ensuring consistency in evaluation across models.

Experimental Setup. We adhere to the training and evaluation protocol established in Dwivedi et al. [29]. For each dataset, we perform hyperparameter tuning via grid search, optimizing the corresponding evaluation metrics: Mean Absolute Error (MAE) for *ZINC-12k*, and classification accuracy for the remaining tasks. We use the AdamW optimizer and train all models for up to 300 epochs, with early stopping based on validation performance. To ensure comparability with prior work, we respect the same parameter budgets used in the original benchmark and maintain the architectural constraints defined for fair evaluation. Each configuration is trained with three random seeds, and we report the average and standard deviation of the results. Hyperparameter ranges used in this set of experiments are summarized in Table 5.

F.4 Long Range Graph Benchmark

Dataset. To evaluate model performance on real-world graphs with significant long-range dependencies, we utilize the *Peptides-func* and *Peptides-struct* benchmarks introduced in Dwivedi et al. [28]. These datasets represent peptide molecules as graphs, where nodes correspond to heavy (non-hydrogen) atoms, and edges denote chemical bonds. *Peptides-func* is a multi-label classification task with 10 functional categories, including antibacterial, antiviral, and signaling-related properties. In contrast, *Peptides-struct* focuses on regression, targeting physical and geometric attributes such as molecular inertia (weighted by atomic mass and valence), atom pair distance extremes, sphericity, and average deviation from a best-fit plane. Together, the two datasets comprise 15,535 peptide graphs and roughly 2.3 million nodes. We adopt the official train/validation/test partitions from Dwivedi et al. [28] and report mean and standard deviation across three different random seeds for each experiment.

Experimental Setup. We follow the evaluation protocol established in Dwivedi et al. [28], including dataset usage, training strategy, and model capacity constraints. Hyperparameter tuning is carried out via grid search, optimizing for Average Precision (AP) in the classification task and Mean Absolute Error (MAE) in the regression task. All models are trained using the AdamW optimizer for up to 300 epochs, with early stopping based on validation performance. To ensure fairness and comparability, all models adhere to the 500K parameter limit, in line with the settings of Dwivedi et al. [28] and Gutteridge et al. [43]. Each configuration is run three times with different weight initializations, and results are averaged. Details of the hyperparameter ranges considered can be found in Table 5.

F.5 Heterophilic Node Classification

747

748

749

750

751

753

754

760

761

762

763

767

768

769

778

779

780

781

782

783

785

787

Dataset. For evaluating performance in heterophilic graph settings, we consider five benchmark tasks introduced by Platonov et al. [72]: Roman-Empire, Amazon-Ratings, Minesweeper, Tolokers, and Questions. These datasets span a diverse range of domains and graph topologies. Roman-Empire is constructed from the Wikipedia article on the Roman Empire, where nodes represent words and edges capture either sequential adjacency or syntactic relations. The task is node classification with 18 syntactic categories, and the underlying graph is sparse and chain-structured, suggesting the presence of long-range dependencies. Amazon-Ratings originates from Amazon's product co-purchasing graph. Nodes correspond to products, linked if they are frequently bought together. The classification task involves predicting discretized average product ratings (five classes), with node features derived from fastText embeddings of product descriptions. Minesweeper is a synthetic dataset modeled as a 100×100 grid. Nodes represent individual cells, with edges connecting adjacent cells. A random 20% of nodes are labeled as mines, and the objective is to classify mine-containing cells based on one-hot features that encode the number of neighboring mines. *Tolokers* is based on the Toloka crowdsourcing platform [59], where each node is a worker (toloker), and edges indicate co-participation on the same project. The task involves binary classification to detect whether a worker has been banned, using node features from user profiles and performance metrics. Questions draws from user interaction data on Yandex Q, a question-answering forum. Nodes represent users, and edges capture answering interactions. The goal is to identify users who remain active, with input features derived from user-provided descriptions. A summary of dataset statistics is provided in Table 4.

Table 4: Statistics of the heterophilic node classification datasets.

	Roman-empire	Amazon-ratings	Minesweeper	Tolokers	Questions
N. nodes	22,662	24,492	10,000	11,758	48,921
N. edges	32,927	93,050	39,402	519,000	153,540
Avg degree	2.91	7.60	7.88	88.28	6.28
Diameter	6,824	46	99	11	16
Node features	300	300	7	10	301
Classes	18	5	2	2	2
Edge homophily	0.05	0.38	0.68	0.59	0.84

Experimental Setup. Our experimental procedure aligns with that of Freitas et al. [35] and Platonov et al. [72]. We conduct a grid search to optimize model performance, using classification accuracy for the *Roman-Empire* and *Amazon-Ratings* tasks, and ROC-AUC for *Minesweeper*, *Tolokers*, and *Questions*. Each model is trained using the AdamW optimizer for a maximum of 300 epochs. Our experiments follow the official dataset splits provided by Platonov et al. [72]. For each model configuration, we perform multiple training runs with different random seeds and report the mean and standard deviation of the results. The hyperparameter grid explored in these experiments is summarized in Table 5.

F.6 Hyperparameters

In Table 5, we summarize the hyperparameter grids used for tuning our TANGO across different benchmarks. Alongside standard training hyperparameters such as learning rate, weight decay, and batch size, our method introduces several additional components. These include the number of unrolled steps L (corresponding to the depth of the energy-based dynamics), the hidden dimension d of node features, and the number of message-passing layers $L_{\rm gnn}$ used within the internal EnergyGNN and TangentGNN modules. In all experiments, we share the architecture depth between EnergyGNN

and TANGENTGNN. We also tune the step size ϵ used in the forward Euler update (Equation (6)), which controls the integration scale of the continuous dynamics. We explore multiple values of L to assess how the number of dynamical steps impacts long-range propagation across different tasks. Details of the complete hyperparameter grid can be found in Table 5.

Table 5: Hyperparameter grids used during model selection for the different benchmark categories: *GraphPropPred* (Diameter, SSSP, Eccentricity), *LRGB* (Peptides-func/struct), *Graph Benchmarks* (ZINC-12k, MNIST, CIFAR-10, CLUSTER, PATTERN), and *Node Classification* (Roman-Empire, Amazon-Ratings, Minesweeper, Tolokers, Questions).

Hyperparameter	GraphPropPred	LRGB	Graph Benchmarks	Node Classification
Unrolled steps L	{1,5,10,20}	{2,4,8,16,32}	{2,4,8,16,32}	{2,4,8,16,32}
GNN layers L_{gnn}	{1,2,4,8,16}	{1,2,4,8,16}	{1,2,4,8,16}	{1,2,4,8,16}
Feature dimension d	{10, 20, 30}	{64, 128,256}	{64, 128, 256}	{64, 128, 256}
Step size ϵ	{0.001, 0.1, 1.0}	{0.001, 0.1, 1.0}	{0.001, 0.1, 1.0}	{0.001, 0.1, 1.0}
Learning rate	{1e-3, 1e-4}	{1e-3, 1e-4}	{1e-3, 1e-4}	{1e-3, 1e-4}
Weight decay	{0,1e-6, 1e-5}	{0, 1e-6, 1e-5}	{0, 1e-6, 1e-5}	{0, 1e-6, 1e-5}
Activation function (σ)	ReLU	ELU, GELU, ReLU	ELU, GELU, ReLU	ELU, GELU, ReLU
Batch size	{32,64,128}	{32,64,128}	{32, 64,128}	N/A

G Additional Results and Comparisons

G.1 Long-Range Benchmark

794

795

796

797

798

799

800

801

802

806

817

We assess the performance of our method on the real-world long-range graph benchmark (LRGB) from [28], focusing on the *Peptides-func* and *Peptides-struct* datasets. We follow the experimental setting in [28], including the 500K parameter budget. All transformer baselines include positional and structural encodings. TANGO does not use additional encodings. The datasets consist of large molecular graphs derived from peptides, where the structure and function of a peptide depend on interactions between distant parts of the graph. Therefore, relying on short-range interactions, such as those captured by local message passing in GNNs, may not be sufficient to excel at this task.

Table 6 provides a comparison of our TANGO model with a wide range of baselines. A broader comparison is presented in Table 10. The results indicate that TANGO outperforms standard MPNNs, transformer-based GNNs, DE-GNNs, and most Multi-hop GNNs.

G.2 Heterophilic Node Classification

We report and compare the performance of our TANGO with other recent benchmarks on the heterophilic node classification datasets from Platonov et al. [72], in Table 7. As can be seen from the Table, TANGO offers strong performance that is similar or better than recent state-of-the-art methods, further demonstrating its effectiveness.

G.3 Additional Comparisons

The comparisons made in Section 3 offer a focused comparison with directly related methods as well as baseline backbones. In addition to that, we now provide a more comprehensive comparison in Table 10 and Table 11, to further facilitate a comprehensive comparison with recent methods. As can be seen, also under these comparisons, our TANGO offers strong performance.

G.4 Ablation Study

Setup. We conduct two key ablation studies to better understand the contributions of the energy function and the tangential flow in TANGO. Specifically, we aim to answer the following questions:

- (i) Does downstream performance benefit from incorporating a tangential term even when the underlying GNN is not the gradient of an energy function?
- (ii) Is the observed improvement due to the tangential nature of the added component, or simply due to additional parameters and network?
- To address these questions, we design two controlled experiments. For comprehensive coverage, we evaluate one representative dataset from each benchmark group: ZINC-12k, Roman-empire,

Table 6: Results for Peptides-func and Peptides-struct (3 training seeds). The **first**, **second**, and **third** best scores are colored.

Model	$\begin{array}{c} \textbf{Peptides-func} \\ \text{AP} \uparrow \end{array}$	Peptides-struct MAE ↓
MPNNs		
GCN [52]	$59.30_{\pm0.23}$	$0.3496_{\pm0.0013}$
GINE [29]	54.98 ± 0.79	$0.3547_{\pm 0.0045}$
GCNII [16]	$55.43_{\pm 0.78}$	$0.3471_{\pm 0.0010}$
GatedGCN [11]	$58.64_{\pm0.77}$	$0.3420_{\pm 0.0013}$
Multi-hop GNNs		
DIGL+MPNN+LapPE [37]	$68.30_{\pm0.26}$	$0.2616_{\pm0.0018}$
MixHop-GCN+LapPE [1]	$68.43_{\pm 0.49}$	$0.2614_{\pm0.0023}$
DRew-GCN+LapPE [43]	71.50 \pm 0.44	0.2536 ± 0.0015
Transformers		
Transformer+LapPE [29]	$63.26_{\pm 1.26}$	$0.2529_{\pm 0.0016}$
SAN+LapPE [54]	$63.84_{\pm 1.21}$	$0.2683_{\pm0.0043}$
GPS+LapPE [75]	$65.35_{\pm0.41}$	$0.2500_{\pm 0.0005}$
DE-GNNs		
GRAND [15]	57.89 ± 0.62	0.3418 ± 0.0015
GraphCON [79]	$60.22_{\pm 0.68}$	$0.2778_{\pm0.0018}$
A-DGN [39]	$59.75_{\pm0.44}$	$0.2874_{\pm0.0021}$
SWAN [40]	67.51 ± 0.39	0.2485 ± 0.0009
PH-DGN [46]	70.12 $_{\pm 0.45}$	0.2465 ± 0.0020
Ours		
TANGOGATEDGCN	68.92 ± 0.40	0.2451 ± 0.0006
TANGOGPS	$70.21_{\pm 0.43}$	$0.2422_{\pm 0.0014}$

Peptides-func, and Diameter. All experiments are run with two backbone architectures, GatedGCN and GPS. For reference, we also report the performance of the original backbones.

Results. For ablation (i), we compare TANGO against a variant we call TANGO-NON-ENERGY, in which the gradient-based energy descent term $\nabla_{\mathbf{H}} V_{\mathcal{G}}(\mathbf{H}^{(\ell)})$ in Equation (6) is replaced by intermediate node features from the same GNN backbone, as detailed in Equation (7). These features are computed using the same architecture but are not guaranteed to correspond to the gradient of any scalar energy function. This setup ensures fairness in capacity while removing the energy-based structure. As shown in Table 8, although both variants benefit from the inclusion of the tangential component, the full TANGO consistently outperforms TANGO-NON-ENERGY, confirming that leveraging a valid energy gradient contributes meaningfully to downstream performance.

For ablation (ii), we isolate the effect of the tangential nature of the added direction. In this variant, denoted Tango-Non-Tangent, we use the same output from the tangential network as in Equation (10) but omit the orthogonal projection step defined in Equation (2). Thus, while we still introduce an additional GNN term into the dynamics, it is not explicitly orthogonal to the energy gradient. Our results in Table 9 show that while this variant improves the performance compared with the baseline backbone, it also results in a drop in performance compared to the full Tango. This highlights the importance of the tangential constraint, and its contribution towards improving the utilization of the learned energy function, as discussed in Section 2.2. Together, these ablations underscore the importance of both components in our design: (i) the principled learned energy descent, and (ii) the structured tangential update, as crucial for effective and flexible feature evolution.

Table 7: Mean test set score and std averaged over the splits from Platonov et al. [72]. **First, second**, and **third** best results for each task are color-coded. We mark each method once – if two variants are among the leading methods, we mark the best-performing variant.

Model	Roman-empire	Amazon-ratings	Minesweeper	Tolokers	Questions
Model	Acc ↑	Acc↑	AUC ↑	AUC ↑	AUC ↑
MPNNs					
GAT	$80.87_{\pm 0.30}$	$49.09_{\pm0.63}$	$92.01_{\pm 0.68}$	$83.70_{\pm0.47}$	$77.43_{\pm 1.20}$
GAT-sep	$88.75_{\pm0.41}$	$52.70_{\pm 0.62}$	$93.91_{\pm 0.35}$	$83.78_{\pm0.43}$	$76.79_{\pm 0.71}$
Gated-GCN	$74.46_{\pm0.54}$	$43.00_{\pm0.32}$	$87.54_{\pm 1.22}$	$77.31_{\pm 1.14}$	$76.61_{\pm 1.13}$
GCN	$73.69_{\pm0.74}$	$48.70_{\pm0.63}$	$89.75_{\pm 0.52}$	$83.64_{\pm0.67}$	$76.09_{\pm 1.27}$
CO - $GNN(\Sigma, \Sigma)$	91.57 $_{\pm 0.32}$	51.28 ± 0.56	$95.09_{\pm 1.18}$	$83.36_{\pm0.89}$	$80.02_{\pm 0.86}$
CO - $GNN(\mu, \mu)$	$91.37_{\pm 0.35}$	$54.17_{\pm 0.37}$	97.31 $_{\pm 0.41}$	$84.45_{\pm 1.17}$	$76.54_{\pm 0.95}$
SAGE	$85.74_{\pm0.67}$	$53.63_{\pm 0.39}$	$93.51_{\pm 0.57}$	$82.43_{\pm0.44}$	$76.44_{\pm0.62}$
Graph Transformers					
Exphormer	89.03 $_{\pm 0.37}$	$53.51_{\pm 0.46}$	$90.74_{\pm 0.53}$	$83.77_{\pm 0.78}$	$73.94_{\pm 1.06}$
NAGphormer	$74.34_{\pm0.77}$	51.26 ± 0.72	$84.19_{\pm 0.66}$	$78.32_{\pm0.95}$	$68.17 {\scriptstyle \pm 1.53}$
GOAT	$71.59_{\pm 1.25}$	$44.61_{\pm 0.50}$	$81.09_{\pm 1.02}$	$83.11_{\pm 1.04}$	$75.76_{\pm 1.66}$
GPS _{GAT+Performer} (RWSE)		$49.92_{\pm 0.68}$	$91.08_{\pm 0.58}$	$84.38_{\pm0.91}$	$77.14_{\pm 1.49}$
GT	$86.51_{\pm 0.73}$	$51.17_{\pm 0.66}$	$91.85_{\pm 0.76}$	$83.23_{\pm0.64}$	$77.95_{\pm0.68}$
GT-sep	$87.32_{\pm0.39}$	$52.18_{\pm0.80}$	$92.29_{\pm 0.47}$	$82.52_{\pm0.92}$	$78.05_{\pm0.93}$
Heterophily-Designated	GNNs				
FAGCN	$65.22_{\pm 0.56}$	$44.12_{\pm0.30}$	$88.17_{\pm 0.73}$	$77.75_{\pm 1.05}$	$77.24_{\pm 1.26}$
FSGNN	$79.92_{\pm 0.56}$	$52.74_{\pm0.83}$	$90.08_{\pm 0.70}$	$82.76_{\pm0.61}$	78.86 \pm 0.92
GBK-GNN	$74.57_{\pm0.47}$	$45.98_{\pm0.71}$	$90.85_{\pm 0.58}$	$81.01_{\pm 0.67}$	$74.47_{\pm0.86}$
GloGNN	$59.63_{\pm0.69}$	$36.89_{\pm0.14}$	$51.08_{\pm 1.23}$	$73.39_{\pm 1.17}$	$65.74_{\pm 1.19}$
GPR-GNN	$64.85_{\pm0.27}$	$44.88_{\pm0.34}$	$86.24_{\pm0.61}$	72.94 ± 0.97	$55.48_{\pm0.91}$
JacobiConv	$71.14_{\pm0.42}$	$43.55_{\pm0.48}$	$89.66_{\pm0.40}$	$68.66_{\pm0.65}$	$73.88_{\pm 1.16}$
Ours					
TANGOGatedGCN	91.89 $_{\pm 0.30}$	$52.60_{\pm 0.53}$	$98.32_{\pm 0.59}$	$85.51 {\scriptstyle \pm 0.98}$	$80.39_{\pm 1.04}$
TANGO _{GPS}	$91.08_{\pm 0.57}$	$53.83_{\pm0.32}$	$98.39_{\pm 0.54}$	$\textbf{85.66}\scriptstyle{\pm 1.01}$	$80.32_{\pm 1.07}$

Table 8: Ablation study on the importance of using a gradient of an energy term in Equation (6).

Model	ZINC-12k MAE↓	Roman-empire Acc. ↑	Peptides-func AP↑	$\begin{array}{c} \textbf{Diameter} \\ \log_{10}(\text{MSE}) \downarrow \end{array}$
GatedGCN TANGO-NON-ENERGY _{GatedGCN} TANGO _{GatedGCN}	$\begin{array}{c} 0.282_{\pm 0.015} \\ 0.138_{\pm 0.014} \\ \textbf{0.128}_{\pm 0.011} \end{array}$	$74.46_{\pm 0.54}\atop86.94_{\pm 0.43}\\\textbf{91.89}_{\pm 0.30}$	$58.64_{\pm 0.77} \ 68.07_{\pm 0.45} \ 68.92_{\pm 0.40}$	$\begin{array}{c} 0.1348_{\pm 0.0397} \\ \text{-}0.5992_{\pm 0.0831} \\ \text{-}\textbf{0.6681}_{\pm 0.0745} \end{array}$
GPS TANGO-NON-ENERGY _{GPS} TANGO _{GPS}	$\begin{array}{c} 0.070_{\pm 0.004} \\ 0.067_{\pm 0.004} \\ \textbf{0.062}_{\pm 0.005} \end{array}$	$87.04_{\pm 0.58} \\ 89.00_{\pm 0.61} \\ \textbf{91.08}_{\pm 0.57}$	$\begin{array}{c} 65.35_{\pm 0.41} \\ 67.58_{\pm 0.39} \\ \textbf{70.21}_{\pm 0.43} \end{array}$	$ \begin{array}{l} \textbf{-0.5121}_{\pm 0.0426} \\ \textbf{-0.7178}_{\pm 0.0729} \\ \textbf{-0.9772}_{\pm 0.0518} \end{array} $

Table 9: The importance of using a tangential term to the energy term in Equation (6).

Model	ZINC-12k MAE↓	Roman-empire Acc. ↑	Peptides-func AP↑	$\begin{array}{c} \textbf{Diameter} \\ \log_{10}(\text{MSE}) \downarrow \end{array}$
GatedGCN TANGO-NON-TANGENT _{GatedGCN} TANGO _{GatedGCN}	$\begin{array}{c} 0.282_{\pm 0.015} \\ 0.186_{\pm 0.016} \\ \textbf{0.128}_{\pm 0.011} \end{array}$	$74.46_{\pm 0.54} \\ 83.59_{\pm 0.48} \\ \textbf{91.89}_{\pm 0.30}$	$58.64_{\pm 0.77} \ 68.01_{\pm 0.52} \ 68.92_{\pm 0.40}$	$\begin{array}{c} 0.1348_{\pm 0.0397} \\ \text{-}0.2193_{\pm 0.0899} \\ \text{-}\textbf{0.6681}_{\pm 0.0745} \end{array}$
GPS TANGO-NON-TANGENT _{GPS} TANGO _{GPS}	$\begin{array}{c} 0.070_{\pm 0.004} \\ 0.066_{\pm 0.010} \\ \textbf{0.062}_{\pm 0.005} \end{array}$	$87.04_{\pm 0.58}$ $88.57_{\pm 0.72}$ 91.08 $_{\pm 0.57}$	$65.35_{\pm 0.41} 67.33_{\pm 0.59} 70.21_{\pm 0.43}$	$ \begin{array}{l} -0.5121 {\pm 0.0426} \\ -0.2916 {\pm 0.0404} \\ \textbf{-0.9772} {\pm 0.0518} \end{array} $

Table 10: Results for Peptides-func and Peptides-struct averaged over 3 training seeds. Baseline results are taken from [28] and [43]. Re-evaluated methods employ the 3-layer MLP readout proposed in [86]. Note that all MPNN-based methods include structural and positional encoding. [‡] means 3-layer MLP readout and residual connections are employed based on [86]. This table is an extended version of the focused Table 6.

Model	Peptides-func	Peptides-struct
Wiodei	AP↑	$MAE\downarrow$
MPNNs		_
GCN	$59.30_{\pm0.23}$	$0.3496_{\pm0.0013}$
GINE	$54.98_{\pm 0.79}$	$0.3547_{\pm 0.0045}$
GCNII	$55.43_{\pm 0.78}$	$0.3471_{\pm 0.0010}$
GatedGCN	$58.64_{\pm0.77}$	$0.3420 {\scriptstyle \pm 0.0013}$
Multi-hop GNNs		
DIGL+MPNN	$64.69_{\pm 0.19}$	$0.3173_{\pm 0.0007}$
DIGL+MPNN+LapPE	$68.30_{\pm0.26}$	$0.2616 {\scriptstyle \pm 0.0018}$
MixHop-GCN	$65.92_{\pm0.36}$	$0.2921_{\pm 0.0023}$
MixHop-GCN+LapPE	$68.43_{\pm0.49}$	$0.2614_{\pm0.0023}$
DRew-GCN	$69.96_{\pm0.76}$	$0.2781_{\pm 0.0028}$
DRew-GCN+LapPE	$71.50_{\pm0.44}$	$0.2536_{\pm0.0015}$
DRew-GIN	$69.40_{\pm 0.74}$	$0.2799_{\pm 0.0016}$
DRew-GIN+LapPE	$71.26_{\pm0.45}$	$0.2606_{\pm0.0014}$
DRew-GatedGCN	$67.33_{\pm 0.94}$	$0.2699_{\pm 0.0018}$
DRew-GatedGCN+LapPE	$69.77_{\pm 0.26}$	$0.2539_{\pm 0.0007}$
Transformers		
Transformer+LapPE	$63.26_{\pm 1.26}$	$0.2529 {\scriptstyle \pm 0.0016}$
SAN+LapPE	$63.84_{\pm 1.21}$	$0.2683_{\pm0.0043}$
GraphGPS+LapPE	$65.35_{\pm0.41}$	$0.2500_{\pm 0.0005}$
Modified and Re-evaluated [‡]		
GCN	$68.60_{\pm 0.50}$	$0.2460_{\pm 0.0007}$
GINE	$66.21_{\pm 0.67}$	$0.2473_{\pm 0.0017}$
GatedGCN	$67.65_{\pm0.47}$	$0.2477_{\pm 0.0009}$
GraphGPS	$65.34_{\pm0.91}$	$0.2509 {\scriptstyle \pm 0.0014}$
DE-GNNs		
GRAND	$57.89_{\pm 0.62}$	$0.3418_{\pm0.0015}$
GraphCON	$60.22_{\pm 0.68}$	$0.2778_{\pm0.0018}$
A-DGN	$59.75_{\pm0.44}$	$0.2874_{\pm 0.0021}$
SWAN	$67.51_{\pm 0.39}$	$0.2485 {\scriptstyle \pm 0.0009}$
Graph SSMs		
Graph-Mamba	$67.39_{\pm0.87}$	$0.2478_{\pm 0.0016}$
GMN	$70.71_{\pm 0.83}$	$0.2473_{\pm 0.0025}$
Ours		
TANGO _{GATEDGCN}	$68.92_{\pm0.40}$	$0.2451_{\pm 0.0006}$
$TANGO_{GPS}$	$70.21_{\pm 0.43}$	$0.2422_{\pm 0.0014}$

Table 11: Mean test set score and std averaged over the splits from Platonov et al. [72]. This table is an extended version of the focused Table 7. Baseline results are reported from [34, 61, 66, 72].

Model	Roman-empire Acc↑	Amazon-ratings Acc ↑	Minesweeper AUC ↑	Tolokers AUC ↑	Questions AUC ↑
MPNNs	1200	1100			
GAT	$80.87_{\pm 0.30}$	$49.09_{\pm0.63}$	$92.01_{\pm 0.68}$	83.70 _{±0.47}	77 43 11 20
GAT-sep	$88.75_{\pm 0.41}$	$52.70_{\pm 0.62}$	$93.91_{\pm 0.35}$	83.70 ± 0.47 83.78 ± 0.43	
GAT (LapPE)	$84.80_{\pm 0.46}$	$44.90_{\pm 0.73}$	$93.50_{\pm 0.54}$	$84.99_{\pm 0.54}$	
GAT (RWSE)	$86.62_{\pm 0.53}$	$48.58_{\pm 0.41}$	$92.53_{\pm 0.65}$	$85.02_{\pm 0.67}$	
GAT (DEG)	$85.51_{\pm 0.56}$	$51.65_{\pm 0.60}$	$93.04_{\pm 0.62}$	$84.22_{\pm 0.81}$	
Gated-GCN	$74.46_{\pm 0.54}$	$43.00_{\pm 0.32}$	$87.54_{\pm 1.22}$	$77.31_{\pm 1.14}$	
GCN GCN	$73.69_{\pm 0.74}$	$48.70_{\pm 0.63}$	$89.75_{\pm 0.52}$	$83.64_{\pm 0.67}$	
GCN (LapPE)	$83.37_{\pm 0.55}$	$44.35_{\pm 0.36}$	$94.26_{\pm 0.49}$	$84.95_{\pm 0.78}$	
GCN (RWSE)	$84.84_{\pm 0.55}$	$46.40_{\pm 0.55}$	$93.84_{\pm 0.48}$	$85.11_{\pm 0.77}$	
GCN (DEG)	$84.21_{\pm 0.47}$	$50.01_{\pm 0.69}$	$94.14_{\pm 0.50}$	$82.51_{\pm 0.83}$	
CO - $GNN(\Sigma, \Sigma)$	$91.57_{\pm 0.32}$	$51.28_{\pm 0.56}$	$95.09_{\pm 1.18}$	$83.36_{\pm 0.89}$	
CO-GNN(μ , μ)	$91.37_{\pm 0.35}$	$54.17_{\pm 0.37}$	$97.31_{\pm 0.41}$	$84.45_{\pm 1.17}$	
SAGE	$85.74_{\pm 0.67}$	$53.63_{\pm 0.39}$	93.51 ± 0.41 93.51 ± 0.57	$82.43_{\pm 0.44}$	
	0017 1 ± 0.07	22.02 ±0.39	70.01±0.37	021.100.44	7011110.02
Graph Transformers	90.02	£2 £1	00.74	92.77	72.04
Exphormer	$89.03_{\pm 0.37}$	$53.51_{\pm 0.46}$	$90.74_{\pm 0.53}$	$83.77_{\pm 0.78}$	
NAGphormer	$74.34_{\pm 0.77}$	$51.26_{\pm 0.72}$	$84.19_{\pm 0.66}$	$78.32_{\pm 0.95}$	
GOAT	$71.59_{\pm 1.25}$	$44.61_{\pm 0.50}$	$81.09_{\pm 1.02}$	$83.11_{\pm 1.04}$	
GPS (LampE)	$82.00_{\pm 0.61}$	$53.10_{\pm 0.42}$	$90.63_{\pm 0.67}$	$83.71_{\pm 0.48}$	
GPS _{GCN+Performer} (LapPE)	$83.96_{\pm 0.53}$	$48.20_{\pm 0.67}$	$93.85_{\pm0.41}$	$84.72_{\pm 0.77}$	
GPS _{GCN+Performer} (RWSE)	84.72 ± 0.65	48.08 ± 0.85	$92.88_{\pm 0.50}$	$84.81_{\pm 0.86}$	
GPS _{GCN+Performer} (DEG)	$83.38_{\pm 0.68}$	$48.93_{\pm 0.47}$	$93.60_{\pm 0.47}$	$80.49_{\pm 0.97}$	
GPS _{GAT+Performer} (LapPE)	$85.93_{\pm 0.52}$	$48.86_{\pm0.38}$	$92.62_{\pm 0.79}$	$84.62_{\pm 0.54}$	
GPS _{GAT+Performer} (RWSE)	$87.04_{\pm 0.58}$	$49.92_{\pm 0.68}$	$91.08_{\pm 0.58}$	$84.38_{\pm 0.91}$	
GPS _{GAT+Performer} (DEG)	$85.54_{\pm 0.58}$	$51.03_{\pm 0.60}$	$91.52_{\pm 0.46}$	$82.45_{\pm 0.89}$	
GPS _{GCN+Transformer} (LapPE)		OOM	$91.82_{\pm 0.41}$	$83.51_{\pm 0.93}$	OOM
GPS _{GCN+Transformer} (RWSE)		OOM	$91.17_{\pm 0.51}$	$83.53_{\pm 1.06}$	OOM
GPS _{GCN+Transformer} (DEG)	OOM	OOM	$91.76_{\pm 0.61}$	$80.82_{\pm 0.95}$	OOM
GPS _{GAT+Transformer} (LapPE)	OOM	OOM	$92.29_{\pm 0.61}$	$84.70_{\pm 0.56}$	OOM
GPS _{GAT+Transformer} (RWSE)	OOM	OOM	$90.82_{\pm 0.56}$	$84.01_{\pm 0.96}$	OOM
GPS _{GAT+Transformer} (DEG)	OOM	OOM	$91.58_{\pm 0.56}$	$81.89_{\pm 0.85}$	OOM
GT	$86.51_{\pm 0.73}$	$51.17_{\pm 0.66}$	$91.85_{\pm 0.76}$	$83.23_{\pm 0.64}$	
GT-sep	$87.32_{\pm0.39}$	$52.18_{\pm0.80}$	$92.29_{\pm 0.47}$	$82.52_{\pm 0.92}$	/8.05±0.9
Heterophily-Designated G					
CPGNN	$63.96_{\pm0.62}$	$39.79_{\pm 0.77}$	$52.03_{\pm 5.46}$	$73.36_{\pm 1.01}$	
FAGCN	$65.22_{\pm 0.56}$	$44.12_{\pm 0.30}$	$88.17_{\pm0.73}$	$77.75_{\pm 1.05}$	
FSGNN	$79.92_{\pm 0.56}$	$52.74_{\pm0.83}$	$90.08_{\pm0.70}$	$82.76_{\pm0.61}$	$78.86_{\pm0.99}$
GBK-GNN	$74.57_{\pm0.47}$	$45.98_{\pm0.71}$	$90.85_{\pm0.58}$	$81.01_{\pm 0.67}$	
GloGNN	$59.63_{\pm 0.69}$	$36.89_{\pm0.14}$	$51.08_{\pm 1.23}$	$73.39_{\pm 1.17}$	
GPR-GNN	$64.85_{\pm0.27}$	$44.88_{\pm0.34}$	$86.24_{\pm0.61}$	$72.94_{\pm 0.97}$	
H2GCN	$60.11_{\pm 0.52}$	$36.47_{\pm0.23}$	$89.71_{\pm 0.31}$	$73.35_{\pm 1.01}$	
JacobiConv	$71.14_{\pm0.42}$	$43.55_{\pm0.48}$	$89.66_{\pm0.40}$	$68.66_{\pm0.65}$	$73.88_{\pm 1.16}$
Graph SSMs					
GMN	$87.69_{\pm0.50}$	$54.07_{\pm0.31}$	$91.01_{\pm 0.23}$	$84.52 _{\pm 0.21}$	_
GPS + Mamba	$83.10_{\pm 0.28}$	$45.13_{\pm 0.97}$	$89.93_{\pm 0.54}$	$83.70_{\pm 1.05}$	_
Ours					
TANGO _{GatedGCN}	$91.89_{\pm0.30}$	$52.60_{\pm0.53}$	$98.32_{\pm 0.59}$	85.51 _{±0.98}	80.39+1.0
TANGOGPS	$91.08_{\pm 0.57}$	$53.83_{\pm 0.32}$	$98.39_{\pm 0.54}$	$85.66_{\pm 1.01}$	