

---

# Gaussian-Process-Based Dynamic Embedding for Textual Networks

---

Pengyu Cheng, Yitong Li, Xinyuan Zhang, Liqun Chen,  
David Carlson, Lawrence Carin  
Duke University  
pengyu.cheng@duke.edu

## Abstract

Textual network embedding aims to learn low-dimensional representations of text-annotated nodes in a graph. Prior works have typically focused on fixed graph structures. However, real-world networks are often dynamic. We address this challenge with a novel end-to-end node-embedding model, called Dynamic Embedding for Textual Networks with a Gaussian Process (DetGP). Because the structure is allowed to be dynamic, our method uses the Gaussian process to take advantage of its non-parametric properties. After training, DetGP can be applied efficiently to dynamic graphs without re-training or backpropagation. To use both local and global graph structure, diffusion is used to model multiple hops between neighbors. The relative importance of global versus local structure for the embeddings is learned automatically. With the non-parametric nature of the Gaussian process, updating the embeddings for a changed graph structure requires only a forward pass through the learned model. Experiments demonstrate the empirical effectiveness of our method compared to baseline approaches, on link prediction and node classification. We further show that DetGP can be straightforwardly and efficiently applied to dynamic textual networks.

## 1 Introduction

Learning latent representations for graph nodes has attracted considerable attention in machine learning, with applications in social networks [5, 17], knowledge bases [25, 32], and recommendation systems [29]. *Textual networks* additionally contain rich semantic information, so text can be included with the graph structure to predict downstream tasks, such as link prediction [12, 31] and node classification [9, 22]. For instance, social networks have links between users, and typically each user has a profile (text). The goal of *textual network embedding* is to learn node embeddings by jointly considering *textual* and *structural* information in the graph.

Most of the aforementioned textual network embedding methods focus on a fixed graph structure [26, 31, 21]. When new network nodes are added to the graph, these frameworks require the whole model be re-trained to update the existing nodes and add representations for the new nodes, leading to high computational complexity. However, networks are often dynamic. For example, in social networks users and relationships between users change over time (*e.g.*, new users, new friends, unfriending, *etc.*). It is impractical to update the full model whenever a new user is added. This paper seeks to address this challenge and learn an embedding method that adapts to a changed graph, without re-training.

Prior dynamic embedding methods usually focus on predicting how the graph structure changes over time, by training on multiple time steps [33, 20, 4]. In such a model, a dynamic network embedding is approximated by multiple steps of fixed network embeddings. In contrast, our method only needs to train on a single graph and can quickly adapt to related graph structures. Additionally, in prior work textual information is rarely included in dynamic graphs. Two exceptions have looked at dynamic network embeddings with changing node attributes [11, 12]. However, both require pre-trained node features, whereas we show that it is more powerful to learn the text encoder in a joint framework with the structural embedding.

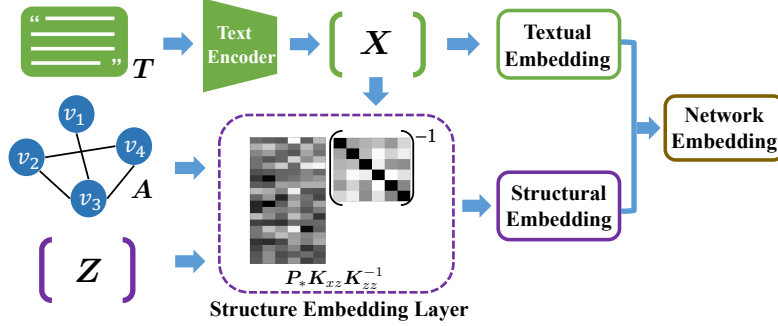


Figure 1: With connection information  $\mathbf{A}$  of the network and textual side information  $\mathbf{T} = \{\mathbf{t}_n\}_{n=1}^N$  as input, DetGP first encodes text  $\mathbf{T}$  to a low-dimensional representation  $\mathbf{X} = \{\mathbf{x}_n\}_{n=1}^N$ , and then infers the structural embeddings by  $\mathbf{X}$  and  $\mathbf{A}$  via a Gaussian process. Inducing points  $\mathbf{Z} = \{\mathbf{z}_m\}_{m=1}^M$  are used to reduce computational complexity. The output network embeddings combine the textual and structural embeddings.

We propose **Dynamic Embedding for Textual Networks with a Gaussian Process (DetGP)**, a novel end-to-end model to perform unsupervised dynamic textual network embedding. DetGP jointly learns *textual* and *structural* features for both fixed and dynamic networks. The textual features indicate the intrinsic attributes of each node based on the text alone, while the structural features reveal the node relationships of the whole community. The structural features utilize both the textual features and the graph topology. This is achieved by smoothing the kernel function in a Gaussian process (GP) with a multi-hop graph transition matrix. This GP-based structure can handle newly added nodes or dynamic edges due to its non-parametric properties [19]. To facilitate fast computation, we learn inducing points to serve as landmarks in the GP [30]. Since the inducing points are fixed after training, computing new embeddings only requires calculating similarity to the inducing points, alleviating computational issues caused by changing graph structure. To evaluate the proposed approach, the learned node embeddings are used for link prediction and node classification. Empirically, DetGP learns improved node representations on several real-world datasets, and outperforms other existing baselines for both static and dynamic networks.

## 2 Model

Assume the input data is given as an undirected graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , where  $\mathcal{V} = \{v_n\}_{n=1}^N$  is the node set and  $\mathcal{E} = \{(v_n, v_{n'}) : v_n, v_{n'} \in \mathcal{V}\}$  is the edge set. Each node  $v_n$  is associated with an  $L_n$ -length text sequence  $\mathbf{t}_n = \{w_i\}_{i=1}^{L_n}$ , where each  $w_i$  is a natural language word. The adjacency matrix  $\mathbf{A} \in \{0, 1\}^{N \times N}$  represents node relationships, where  $A_{nn'} = 1$  if  $(v_n, v_{n'}) \in \mathcal{E}$  and  $A_{nn'} = 0$  otherwise. Our objective is to learn a low-dimensional embedding vector  $\mathbf{h}_n$  for each node  $v_n \in \mathcal{V}$ , that captures both textual and structural features of graph  $\mathcal{G}$ .

Figure 1 gives the framework of the proposed model, DetGP. Text  $\mathbf{t}_n$  is input to a text encoder  $\mathbf{g}_\theta(\cdot)$  with parameters  $\theta$ ; in the Supplementary Material we describe further details. The output  $\mathbf{x}_n = \mathbf{g}_\theta(\mathbf{t}_n)$  is the *textual embedding* of node  $v_n$ . This textual embedding is both part of the complete embedding  $\mathbf{h}_n$  and an input into the structural embedding layer (dotted purple box) that is combined with the graph structure in a GP framework, discussed in Section 2.1. In addition, multiple hops are modeled in this embedding layer to better reflect the graph architecture and use both local and global graph structure. To scale up the model to large datasets, we adopt the idea of inducing points [24, 23], which serve as nonuniformly situated grid points in the model. The output structural embeddings are denoted as  $\mathbf{s}_n$ , which are combined to form the complete node embedding  $\mathbf{h}_n = [\mathbf{x}_n; \mathbf{s}_n]$ .

The model is trained by using the negative sampling loss [26], where neighbor nodes should be more similar than non-neighbor nodes (described in Section 2.2). When the graph structure changes, the node embeddings are updated by a single forward-propagation step without relearning any model parameters. This property comes from the non-parametric nature of the GP-based structure, and it greatly increases computational efficiency for dynamic graphs.

### 2.1 Structural Embedding Layer

The structural embedding layer transforms the encoded text feature  $\mathbf{x}$  to structural embedding  $\mathbf{s}$  using a GP in conjunction with the graph topology. Before introducing the GP, we introduce the multi-hop transition matrix  $\mathbf{P}_*$  that will smooth the GP kernel.

**Multi-hop transition matrix:** Suppose  $\mathbf{P}$  is the normalized transition matrix, *i.e.*, a normalized version of  $\mathbf{A}$  where each row sums to one and  $P_{nn'}$  represents the probability of a transition from node  $n$  to node  $n'$ . If  $\mathbf{P}$  represents the transition from a single hop, then higher orders of  $\mathbf{P}$  will give multi-hop transition probabilities. Specifically,  $\mathbf{P}^j$  is the  $j$ th power of  $\mathbf{P}$ , where  $P_{nn'}^j$  gives the probability of transitioning from node  $n$  to node  $n'$  after  $j$  random hops on the graph. Different powers of  $\mathbf{P}$  provide different levels of smoothing on the graph, and vary from using local ( $\mathbf{P}^0$  or  $\mathbf{P}^1$ ) and global (large powers of  $\mathbf{P}$ ) graph structure for the node embeddings. Equation (1) can be viewed as a generalized form for DeepWalk [17] or GloVe [16]. In practice, learning the weights is more robust than hand-engineering them [1].

$$\mathbf{P}_* = \sum_{j=0}^J \alpha_j \mathbf{P}^j, \quad \text{s.t. } \sum_{j=0}^J \alpha_j = 1, \quad (1)$$

where  $J$  is the maximum number of steps considered, and  $\alpha$  are the learnable weights. The constraint that  $\sum_{j=0}^J \alpha_j = 1$  in (1) is implemented by a softmax function. Note that  $\mathbf{P}^0 = \mathbf{I}_N$  is an identity matrix, which treats each node independently. In contrast, a large power of  $\mathbf{P}$  would typically be very smooth after taking many hops. Therefore,  $\alpha$  can learn the importance of local ( $\mathbf{P}^0$  or  $\mathbf{P}^1$ ) and global (large powers of  $\mathbf{P}$ ) graph structure for the node embeddings. Equation (1) can be viewed as a generalized form for DeepWalk [17] or GloVe [16]. In practice, learning the weights is more robust than hand-engineering them [1].

**GP structural embedding prior:** We define a latent function  $f(\mathbf{x})$  over the textual embedding  $\mathbf{x}$  with a GP prior  $f(\mathbf{x}) \sim \mathcal{GP}(\mathbf{0}, k(\mathbf{x}_n, \mathbf{x}_{n'}))$ . Inspired by [15], instead of using this GP directly to determine the embedding, the learned graph diffusion is used on top of this Gaussian process. For finite samples, the combination of the graph diffusion and the GP yields a conditional structural embedding that can be expressed as a multivariate Gaussian distribution:

$$p([s_{1i}, \dots, s_{Ni}]^\top | [\mathbf{x}_1, \dots, \mathbf{x}_N]) = \mathcal{N}(\mathbf{0}, \mathbf{P}_*^\top \mathbf{K}_{XX} \mathbf{P}_*), \quad (2)$$

where  $[\mathbf{K}_{XX}]_{nn'} = k(\mathbf{x}_n, \mathbf{x}_{n'})$  and  $i$  is a index of our structure embedding feature. Each dimension of the structural embedding follows this Gaussian distribution with the same covariance matrix  $\mathbf{K}_{XX}$  smoothed by  $\mathbf{P}_*$ . In the Supplementary Material, we discuss the selection of different kernels.

**Inducing Points:** GP models are well-known to suffer from computational complexity with large data size  $N$ . To scale up the model, we use the inducing points based on the variational Gaussian process [24]. Denote inducing points as  $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_M]^\top$  with  $M < N$ , and their corresponding learnable embeddings as  $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_M]^\top$ , which follow the same GP function. The textual and structural embeddings of real data samples are denoted as  $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_N]^\top$  and  $\mathbf{S} = [s_{1i}, \dots, s_{Ni}]^\top$ . Given inducing points, the conditional distribution of our structural embeddings is

$$p(\mathbf{S}_i | \mathbf{X}, \mathbf{Z}, \mathbf{U}) = \mathcal{N}(\boldsymbol{\mu}_{\mathbf{S}_i | \mathbf{Z}}, \boldsymbol{\Sigma}_{\mathbf{S} | \mathbf{Z}}), \quad (3)$$

$$\text{where } \boldsymbol{\mu}_{\mathbf{S}_i | \mathbf{Z}} = \mathbf{P}_*^\top \mathbf{K}_{XZ} (\mathbf{K}_{ZZ} + \sigma \mathbf{I}_M)^{-1} \mathbf{U}_i,$$

$$\boldsymbol{\Sigma}_{\mathbf{S} | \mathbf{Z}} = \mathbf{P}_*^\top \mathbf{K}_{XX} \mathbf{P}_* - \mathbf{P}_*^\top \mathbf{K}_{XZ} (\mathbf{K}_{ZZ} + \sigma \mathbf{I}_M)^{-1} \mathbf{K}_{ZX} \mathbf{P}_*.$$

Here  $[\mathbf{K}_{XZ}]_{nm} = k(\mathbf{x}_n, \mathbf{z}_m)$  and  $[\mathbf{K}_{ZZ}]_{mm'} = k(\mathbf{z}_m, \mathbf{z}_{m'})$ . The subscript  $i$  indicates the  $i$ th column of a matrix ( $\mathbf{S}_i$  is the concatenation of the  $i$ th element from all node structural embeddings). Each dimension of  $\mathbf{S}$  has a multivariate Gaussian distribution with unique mean value  $\boldsymbol{\mu}_{\mathbf{S}_i | \mathbf{Z}}$  but the same covariance  $\boldsymbol{\Sigma}_{\mathbf{S} | \mathbf{Z}}$ . Theoretically, we can give a posterior of  $\mathbf{u}$  and get the marginal distribution of  $\mathbf{s}$  by integrating  $\mathbf{u}$  out. However, the integral on  $\mathbf{u}$  does not have a closed form. As an approximation, we use a deterministic function  $\mathbf{S} = \mathbf{P}_*^\top \mathbf{K}_{XZ} (\mathbf{K}_{ZZ} + \sigma \mathbf{I}_M)^{-1} \mathbf{U}$  for the learned structural embedding.

## 2.2 Algorithm Outline

The structural embedding  $\mathbf{s}_n$  and the textual embedding  $\mathbf{x}_n$  are concatenated to form the final node embedding  $\mathbf{h}_n = [\mathbf{x}_n; \mathbf{s}_n]$ . To learn the embeddings in an unsupervised manner, existing works adopt the technique of negative sampling [26], which tries to maximize the conditional probability of a nodal embedding given its neighbors, while maintaining a low conditional probability for non-neighbors. In the proposed framework, this loss is

$$\mathcal{L} = -\frac{1}{|\mathcal{E}|} \sum_{(v_n, v_{n'}) \in \mathcal{E}} \log(\sigma(\mathbf{h}_n^\top \mathbf{h}_{n'})) - \frac{1}{N_s} \sum_{(v_n, v_{n'}) \notin \mathcal{E}} \log[1 - \sigma(\mathbf{h}_n^\top \mathbf{h}_{n'})], \quad (4)$$

where  $N_s = \#\{(v_n, v_{n'}) \notin \mathcal{E}\}$  is a weighting constant. Equation (4) maximizes the inner product among neighbors in the graph while minimizing the similarity among non-neighbors. Our model is trained end-to-end by taking gradients of loss  $\mathcal{L}$  with respect to  $\boldsymbol{\theta}$ ,  $\mathbf{Z}$  and  $\mathbf{U}$ . The inducing points are initialized as the k-means centers of the encoded text features. Then,  $\mathbf{Z}$  and the text encoder are jointly trained to minimize the loss function.

For newly introduced nodes  $\mathcal{V}_{new}$  with text  $\mathbf{T}_{new}$ , the transition matrix  $\mathbf{P}_{new}$  is first updated, and the embeddings are obtained directly without additional back-propagation. Specifically, we first compute  $\mathbf{X}_{new} = \mathbf{g}_\theta(\mathbf{T}_{new})$  from the text encoder. Then with  $\mathbf{X}' = [\mathbf{X}, \mathbf{X}_{new}]$ , the structural embedding of all nodes can be computed as  $[\mathbf{S}, \mathbf{S}_{new}] = \sum_{j=0}^J \alpha_j \mathbf{P}_{new}^j \mathbf{K}_{X'Z} (\mathbf{K}_{ZZ} + \sigma \mathbf{I}_M)^{-1} \mathbf{U}$ .

### 3 Experiments

To demonstrate the efficacy of DetGP embeddings, we conduct experiments on both static and dynamic textual networks. Here we mainly focus on analyzing results on graphs with newly added nodes. The results on static networks and the experiment setup details are shown in the Supplementary Material. Here we use word embedding average (Wavg) [27] as our text encoder.

% Training Nodes	Cora				HepTh			
	10%	30%	50%	70%	10%	30%	50%	70%
Only Text (Wavg)	61.2	77.9	87.9	90.3	68.3	83.7	84.2	86.9
Neighbor-Aggregate (Max-Pooling)	54.6	69.1	78.7	87.3	59.6	78.3	79.9	80.7
Neighbor-Aggregate (Mean)	61.8	78.4	88.0	91.2	68.2	83.9	85.5	88.3
GraphSAGE (Max-Pooling)	62.1	78.6	88.6	92.4	68.4	85.8	88.1	91.2
GraphSAGE (Mean)	62.2	79.1	88.9	92.6	69.1	85.9	89.0	92.4
DetGP	<b>62.9</b>	<b>81.1</b>	<b>90.9</b>	<b>93.0</b>	<b>70.7</b>	<b>86.6</b>	<b>90.7</b>	<b>93.3</b>

Table 1: Test AUC scores for link prediction on *Cora* and *HepTh* datasets.

% Training Nodes	Cora				DBLP			
	10%	30%	50%	70%	10%	30%	50%	70%
Only Text (Wavg)	60.2	76.3	83.5	84.8	56.7	67.9	70.4	73.5
Neighbor-Aggregate (Max-Pooling)	55.8	70.2	78.4	80.5	51.8	60.5	68.3	70.6
Neighbor-Aggregate (Mean)	60.1	77.2	84.1	85.0	56.8	68.2	71.3	74.7
GraphSAGE (Max-Pooling)	61.3	78.2	85.1	86.3	58.9	69.1	72.4	74.9
GraphSAGE (Mean)	61.4	78.4	85.5	<b>86.6</b>	59.0	69.3	72.7	75.1
DetGP	<b>62.1</b>	<b>79.3</b>	<b>85.8</b>	<b>86.6</b>	<b>60.2</b>	<b>70.1</b>	<b>73.2</b>	<b>75.8</b>

Table 2: Test Macro-F1 scores for multi-label node classification on *Cora* and *DBLP* dataset.

Previous works [26, 31, 21] on textual network embedding require the overall connection information to train the structural embedding, which cannot directly assign (without re-training) structural embeddings to a new coming node with connection information unknown during training. Therefore, the aforementioned methods cannot be applied to dynamic networks. To obtain comparable baselines to DetGP, we propose two strategies, based on the idea of (GraphSAGE) [7]: **Neighbor-Aggregate** and **GraphSAGE**. Details about the strategies are shown in the Supplementary Material.

We evaluate the dynamic embeddings for test nodes on link prediction and node classification tasks. For both tasks, we split the nodes into training and testing sets with different proportions (10%, 30%, 50%, 70%). When embedding new testing nodes, only their textual attributes and connections with existing training nodes are provided. For the link prediction, we predict the edges between testing nodes based on the inner product between their node embeddings; for node classification, an SVM classifier is trained based on embeddings of training nodes. When new nodes come, we first embed the nodes using the trained model and then use the pre-learned SVM to predict their labels.

The results of link prediction and node classification are given in Tables 1 and 2, respectively. The proposed DetGP significantly outperforms other baselines, especially when the proportion of training set is small. A reasonable explanation is, when the training set is small, new nodes will have few connections with the training nodes, which causes high variance in the results of aggregating neighborhood embeddings. However, instead of aggregating, the proposed DetGP infers the structural embedding via a Gaussian process with pre-learned inducing points, which is more robust than the information passed by neighbor nodes.

### 4 Conclusions

We propose a novel textual network embedding framework that learns representative node embeddings for static textual network, and also effectively adapts to dynamic graph structures. This is achieved by introducing a GP network structural embedding layer, which first maps each node to the inducing points, and then embeds them by taking advantage of the non-parametric representation. We also consider multiple hops to weight local and global graph structures. The graph structure is injected in the kernel matrix, where the kernel between two nodes use the whole graph information based on multiple hops. Our final embedding contains both *structural* and *textual* information. Empirical results demonstrate the practical effectiveness of the proposed algorithm.

## References

- [1] Abu-El-Haija, S.; Perozzi, B.; Al-Rfou, R.; and Alemi, A. A. 2018. Watch your step: Learning node embeddings via graph attention. In *NeurIPS*, 9180–9190.
- [2] Airoldi, E. M.; Blei, D. M.; Fienberg, S. E.; and Xing, E. P. 2008. Mixed membership stochastic blockmodels. *JMLR* 1981–2014.
- [3] Cer, D.; Yang, Y.; Kong, S.-y.; Hua, N.; Limtiaco, N.; John, R. S.; Constant, N.; Guajardo-Cespedes, M.; Yuan, S.; Tar, C.; et al. 2018. Universal sentence encoder. *arXiv preprint arXiv:1803.11175*.
- [4] Du, L.; Wang, Y.; Song, G.; Lu, Z.; and Wang, J. 2018. Dynamic network embedding: An extended approach for skip-gram based network embedding. In *IJCAI*, 2086–2092.
- [5] Fan, W.; Ma, Y.; Li, Q.; He, Y.; Zhao, E.; Tang, J.; and Yin, D. 2019. Graph neural networks for social recommendation. *arXiv preprint arXiv:1902.07243*.
- [6] Grover, A., and Leskovec, J. 2016. node2vec: Scalable feature learning for networks. In *SIGKDD*, 855–864. ACM.
- [7] Hamilton, W.; Ying, Z.; and Leskovec, J. 2017. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, 1024–1034.
- [8] Hanley, J. A., and McNeil, B. J. 1982. The meaning and use of the area under a receiver operating characteristic (roc) curve. *Radiology* 143(1):29–36.
- [9] Kipf, T. N., and Welling, M. 2017. Semi-supervised classification with graph convolutional networks. *ICLR*.
- [10] Leskovec, J.; Kleinberg, J.; and Faloutsos, C. 2007. Graph evolution: Densification and shrinking diameters. *TKDD*.
- [11] Li, J.; Dani, H.; Hu, X.; Tang, J.; Chang, Y.; and Liu, H. 2017. Attributed network embedding for learning in a dynamic environment. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, 387–396. ACM.
- [12] Li, J.; Cheng, K.; Wu, L.; and Liu, H. 2018. Streaming link prediction on dynamic attributed networks. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, 369–377. ACM.
- [13] Lin, Z.; Feng, M.; Santos, C. N. d.; Yu, M.; Xiang, B.; Zhou, B.; and Bengio, Y. 2017. A structured self-attentive sentence embedding. *ICLR*.
- [14] Melamud, O.; Goldberger, J.; and Dagan, I. 2016. context2vec: Learning generic context embedding with bidirectional lstm. In *SIGNLL*, 51–61.
- [15] Ng, Y. C.; Colombo, N.; and Silva, R. 2018. Bayesian semi-supervised learning with graph gaussian processes. In *NeurIPS*.
- [16] Pennington, J.; Socher, R.; and Manning, C. 2014. Glove: Global vectors for word representation. In *EMNLP*, 1532–1543.
- [17] Perozzi, B.; Al-Rfou, R.; and Skiena, S. 2014. Deepwalk: Online learning of social representations. In *SIGKDD*, 701–710. ACM.
- [18] Powers, D. M. 2011. Evaluation: from precision, recall and f-measure to roc, informedness, markedness and correlation.
- [19] Rasmussen, C. E. 2006. Gaussian processes for machine learning. In *Springer*.
- [20] Seo, Y.; Defferrard, M.; Vandergheynst, P.; and Bresson, X. 2018. Structured sequence modeling with graph convolutional recurrent networks. In *International Conference on Neural Information Processing*, 362–373. Springer.

- [21] Shen, D.; Zhang, X.; Henaio, R.; and Carin, L. 2019. Improved semantic-aware network embedding with fine-grained word alignment. *EMNLP*.
- [22] Tang, J.; Qu, M.; Wang, M.; Zhang, M.; Yan, J.; and Mei, Q. 2015. Line: Large-scale information network embedding. In *Proceedings of the 24th international conference on world wide web*, 1067–1077. International World Wide Web Conferences Steering Committee.
- [23] Titsias, M., and Lawrence, N. D. 2010. Bayesian gaussian process latent variable model. In *AISTATS*, 844–851.
- [24] Titsias, M. 2009. Variational learning of inducing variables in sparse gaussian processes. In *AISTATS*, 567–574.
- [25] Trivedi, R.; Dai, H.; Wang, Y.; and Song, L. 2017. Know-evolve: Deep temporal reasoning for dynamic knowledge graphs. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 3462–3471. JMLR. org.
- [26] Tu, C.; Liu, H.; Liu, Z.; and Sun, M. 2017. Cane: Context-aware network embedding for relation modeling. In *ACL*.
- [27] Wieting, J.; Bansal, M.; Gimpel, K.; and Livescu, K. 2015. Towards universal paraphrastic sentence embeddings. *arXiv preprint arXiv:1511.08198*.
- [28] Yang, C.; Liu, Z.; Zhao, D.; Sun, M.; and Chang, E. 2015. Network representation learning with rich text information. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.
- [29] Ying, R.; He, R.; Chen, K.; Eksombatchai, P.; Hamilton, W. L.; and Leskovec, J. 2018. Graph convolutional neural networks for web-scale recommender systems. In *SIGKDD*, 974–983. ACM.
- [30] Yu, K., and Chu, W. 2008. Gaussian process models for link analysis and transfer learning. In *NIPS*, 1657–1664.
- [31] Zhang, X.; Li, Y.; Shen, D.; and Carin, L. 2018a. Diffusion maps for textual network embedding. In *NeurIPS*.
- [32] Zhang, Y.; Yao, Q.; Shao, Y.; and Chen, L. 2018b. Nscaching: Simple and efficient negative sampling for knowledge graph embedding. *arXiv preprint arXiv:1812.06410*.
- [33] Zhou, L.; Yang, Y.; Ren, X.; Wu, F.; and Zhuang, Y. 2018. Dynamic network embedding by modeling triadic closure process. In *Thirty-Second AAAI Conference on Artificial Intelligence*.

## A Model Details

### A.1 Text Encoder

There are many existing text encoders [14, 3, 13], often based on deep neural networks. However, using a deep neural network encoder can overfit on graphs because of the relatively small size of textual data [31]. Therefore, various encoders are proposed to extract rich textual information specifically from graphs [26, 31, 21]. In general, we aim to learn a text encoder  $\mathbf{g}_\theta$  with parameters  $\theta$  that encodes semantic features  $\mathbf{x}_n = \mathbf{g}_\theta(\mathbf{t}_n)$ . A simple and effective text encoder is the word embedding average (Wavg)  $\mathbf{g}_\theta(\mathbf{t}) = \frac{1}{L_n} \sum_{i=1}^{L_n} \mathbf{v}_i$ , where  $\mathbf{v}_i$  is the corresponding embedding of  $w_i$  from the sequence  $\mathbf{t}_n = \{w_i\}_{i=1}^{L_n}$ . This is implemented by a learnable look-up table. [31] proposed a diffused word average encoder (DWavg) to leverage textual information over multiple hops on the network. Because DetGP focuses mainly on the structural embeddings, we do not focus on developing a new text encoder. Instead, we show that DetGP has compatibility with different text encoders, and our experiments use these two text encoders (Wavg and DWavg).

### A.2 Kernel Selection

While there are many ways to define the kernel function, excessive non-linearity is not required or desired in this embedding layer because the text encoder is highly non-linear and high-dimensional. In practice, the first-degree polynomial kernel, written as

$$k(\mathbf{x}_n, \mathbf{x}_{n'}) = \mathbf{x}_n^\top \mathbf{x}_{n'} + C, \quad C > 0 \quad (5)$$

outperforms others due to its numerical stability. Empirically, the linear kernel in Eq. (5) speeds up computation and increases model stability.

### A.3 Analysis of the Structural Embedding

We analyze the kernel function in Eq. (2) and (5) to show how the graph structure is used in the embedding layer. Denote  $P_{nn'}^j$  (Eq. (1)) as the transition probability from node  $n$  to node  $n'$  in  $j$  hops, then the correlation between node  $n$  and  $n'$  in Eq. (2) can be expanded as

$$\begin{aligned} \text{cov}(s_{ni}, s_{n'i}) &= \alpha_0^2 k(\mathbf{x}_n, \mathbf{x}_{n'}) + \alpha_0 \sum_{j=1}^J \alpha_j \sum_{r=1}^N P_{nr}^j k(\mathbf{x}_n, \mathbf{x}_r) \\ &+ \alpha_0 \sum_{j=1}^J \alpha_j \sum_{r=1}^N P_{n'r}^j k(\mathbf{x}_{n'}, \mathbf{x}_r) \\ &+ \sum_{j=1}^J \sum_{j'=1}^J \sum_{\substack{r=1 \\ r \neq n, n'}}^N \sum_{\substack{r'=1 \\ r' \neq n, n'}}^N \alpha_j \alpha_{j'} P_{nr}^j P_{n'r'}^{j'} k(\mathbf{x}_{n'}, \mathbf{x}_{r'}) \end{aligned} \quad (6)$$

The covariance is the same for all indices  $i$ . The first term in Eq. (6) measures the kernel function between  $\mathbf{x}_n$  and  $\mathbf{x}_{n'}$ . The next two terms show the relationship between  $\mathbf{x}_n$  and the weighted multi-hop neighbors of  $\mathbf{x}_{n'}$  and vice versa.  $\alpha$  controls how much different hops are used. The last term is the pairwise-weighted higher order relationship between any two nodes in the graph, except  $n$  and  $n'$ . The covariance structure uses the whole graph and learns how to balance local and global information. If node  $n$  has no edges, then it will not be influenced by other nodes besides textual similarity. In contrast, a node with dense edge connections will be smoothed by its neighbors.

With the inducing points, Equation (6) can be modified as  $\text{cov}(s_{ni}, u_{mi}) = \alpha_0 k(\mathbf{x}_n, \mathbf{z}_m) + \sum_{j=1}^J \alpha_j \sum_{n'=1}^N P_{nn'}^j k(\mathbf{x}_{n'}, \mathbf{z}_m)$ . The covariance between node  $v_n$  and the inducing points includes the local information  $k(\mathbf{x}_n, \mathbf{z}_m)$ , as well as the smoothed effect from  $\sum_{j=1}^J \alpha_j \sum_{n'=1}^N P_{nn'}^j k(\mathbf{x}_{n'}, \mathbf{z}_m)$ . This can also be viewed as feature smoothing over neighbors. Since inducing points do not contain links to other inducing points, there is no smoothing function for them. Each inducing point can be viewed as a node that already includes global graph information.

## B Experiments

### B.1 Setup Details

For a fair comparison with previous work, we follow the setup in [26, 31, 21], where the embedding for each node has dimension 200, a concatenation of a 100-dimensional textual embedding and a 100-dimensional structural embedding. We evaluate our DetGP base on two text encoders: the word embedding average (Wavg) encoder and the diffused word embedding average (DWavg) encoder from Zhang et al. [31], introduced in Section A.1. The maximum number of hops  $J$  in  $P^*$  is set to 3.

For static network embedding, we follow setups in [26, 31, 21]. In the following, we evaluate the graph embeddings on link prediction and node classification on the following real-world datasets:

- **Cora** is a paper citation network, with a total of 2,277 vertices and 5,214 edges in the graph, where only nodes with text are kept. Each node has a text abstract about machine learning and belongs to one of seven categories.
- **DBLP** is a paper citation network with 60,744 nodes and 52,890 edges. Each node represents one paper in computer science in one of four categories: database, data mining, artificial intelligence, and computer vision.
- **HepTh** (High Energy Physics Theory) [10] is another paper citation network. The original dataset contains 9,877 nodes and 25,998 edges. We only keep nodes with associated text, so this is limited to 1,038 nodes and 1,990 edges.

### B.2 Embedding Assigning Strategies

The two embedding assigning strategies are: (a) **Neighbor-Aggregate**: aggregating the *structural embeddings* from the neighbors in the training set, as the structural embedding for the new node; (b) **GraphSAGE**: aggregating the *textual embeddings* from the neighbors, then passing through a fully-connected layer to get the new node’s structural embedding. For neighborhood information aggregating, we use the mean aggregator and the max-pooling aggregator as mentioned in [7].

In both tasks, the Neighbor-Aggregate strategy with mean aggregator shows slight improvement to the baseline with only a text encoder. However, it does not work well with the max-pooling aggregator, implying that the unsupervised max-pooling on pre-trained neighbor structural embeddings cannot learn a good representation. The GraphSAGE strategies (with both mean and pooling aggregator) show notable improvements compared with Wavg and Neighbor-Aggregate. Unlike the unsupervised pooling, the GraphSAGE pooling aggregator is trained with a fully-connected layer on top, which shows comparable result to the mean aggregator.

### B.3 Link Prediction

The link prediction task seeks to infer if two nodes are connected, based on the learned embeddings. This standard task tests if the embedded node features contain graph connection information. For a given network, we randomly keep a certain percentage (15%, 35%, 55%, 75%, 95%) of edges and learn embeddings. At test time, we calculate the inner product of pairwise node embedding. A large inner product value indicates a potential edge between two nodes. The AUC score [8] is computed in this setting to evaluate the performance. The results are shown in Table 3 on Cora and HepTh. Since the DBLP dataset only has 52,890 edges which is far too sparse compared with the node number 60,744, we do not evaluate the AUC score on it as a consequence of high variance from sampling edges. The first four models only embed structural features, while the remaining alternatives use both textual and structural embeddings. We also provide the DetGP results of with only textual embeddings and only structure embeddings for ablation study.

From Table 3, adding textual information in the embedding can improve the link prediction result by a large margin. Even using only textual embeddings, DetGP gains significant improvement compared with only structure-based methods, and achieves competitive performance compared with other text-based embedding methods. Using only structural information is slightly better than using only textual embeddings, since link prediction is a more structure-dependent task, which also indicates that DetGP learns inducing points  $Z$  that can effectively represent the network structure. Compared with other textual network embedding methods, DetGP has very competitive AUC scores, especially when only given a small percentage of edges. Noting that for our methods the text encoders come



% Training Edges	Cora					HepTh				
	15%	35%	55%	75%	95%	15%	35%	55%	75%	95%
MMB [2]	54.7	59.5	64.9	71.1	75.9	54.6	57.3	66.2	73.6	80.3
node2vec [6]	55.9	66.1	78.7	85.9	88.2	57.1	69.9	84.3	88.4	89.2
LINE [22]	55.0	66.4	77.6	85.6	89.3	53.7	66.5	78.5	87.5	87.6
DeepWalk [17]	56.0	70.2	80.1	85.3	90.3	55.2	70.0	81.3	87.6	88.0
TADW [28]	86.6	90.2	90.0	91.0	92.7	87.0	91.8	91.1	93.5	91.7
CANE [26]	86.8	92.2	94.6	95.6	97.7	90.0	92.0	94.2	95.4	96.3
DMATE [31]	91.3	93.7	96.0	97.4	98.8	NA	NA	NA	NA	NA
WANE [21]	91.7	94.1	96.2	<b>97.5</b>	<b>99.1</b>	92.3	95.7	97.5	97.7	<b>98.7</b>
DetGP (Wavg) only Text	83.4	89.1	89.9	90.9	92.3	86.5	89.6	90.2	91.5	92.6
DetGP (Wavg) only Struct	85.4	89.7	91.0	92.7	94.1	89.7	92.1	93.5	94.8	95.1
DetGP (Wavg)	92.8	94.8	95.5	96.2	97.5	93.2	95.1	97.0	97.3	97.9
DetGP (DWavg)	<b>93.4</b>	<b>95.2</b>	<b>96.3</b>	<b>97.5</b>	98.8	<b>94.3</b>	<b>96.2</b>	<b>97.7</b>	<b>98.1</b>	98.5

Table 3: AUC scores for link prediction on the *Cora* and *HepTh* dataset.

% Training Nodes	Cora				DBLP			
	10%	30%	50%	70%	10%	30%	50%	70%
LINE[22]	53.9	56.7	58.8	60.1	42.7	43.8	43.8	43.9
TADW [28]	71.0	71.4	75.9	77.2	67.6	68.9	69.2	69.5
CANE [26]	81.6	82.8	85.2	86.3	71.8	73.6	74.7	75.2
DMTE [31]	81.8	83.9	86.3	87.9	72.9	74.3	75.5	76.1
WANE [21]	81.9	83.9	86.4	88.1	NA	NA	NA	NA
DetGP (Wavg) only Text	78.1	81.2	84.7	85.3	71.4	73.3	74.2	74.9
DetGP (Wavg) only Struct	70.9	79.7	81.5	82.3	70.0	71.4	72.6	73.3
DetGP (Wavg)	80.5	85.4	86.7	88.5	76.9	78.3	79.1	79.3
DetGP (DWavg)	<b>83.1</b>	<b>87.2</b>	<b>88.2</b>	<b>89.8</b>	<b>78.0</b>	<b>79.3</b>	<b>79.6</b>	<b>79.8</b>

Table 4: Test Macro-F1 scores for multi-label node classification on *Cora* and *HepTh* dataset.

from the baselines Wavg and DWavg [31], the performance gain should come from the proposed structural embedding framework.

#### B.4 Node Classification

Node classification requires high-quality *textual embeddings* because *structural embeddings* alone do not accurately reflect node category. Therefore, we only compare to methods designed for textual network embedding. After training converges, a linear SVM classifier is learned on the trained node embeddings and performance is estimated by a hold-out set. In Table 4, we compare our methods (Wavg+DetGP, DWavg+DetGP) with recent textual network embedding methods under different proportions (10%, 30%, 50%, 70%) of given nodes in the training set. Following the setup in Zhang et al. [31], the evaluation metric is Macro-F1 score [18]. We test on the Cora and DBLP datasets, which have group label information, where DetGP yields the best performance under all situations. This demonstrates that the proposed model can learn both representative textual and structural embeddings. The ablation study results (only textual embeddings vs. only structural embeddings) demonstrates that textual attributes are more important than edge connections in classification task. To describe the effect of learning the weighting in the diffusion, for the experiment on Cora with 10% nodes given for training, the learned weights in  $\alpha$  are [0.58, 0.12, 0.24, 0.05]. Thus, local and second order transition features are more important.

#### B.5 Inducing Points

Figure 2 gives the t-SNE visualization of the learned DetGP structural embeddings on the Cora citation dataset. The model is learned using all edges and all of the nodes with their textual information. We set the number of inducing points to  $M = 20$ . To avoid the computational instability caused by the inverse matrix  $K_{ZZ}^{-1}$ , we update inducing points with a smaller learning rate, which is set to one-tenth of the learning rate for the text encoder. The inducing points  $z$  are visualized as red filled circles in Figure 2. Textual embeddings are plotted with 7 different colors, representing the 7 node classes. Note that the inducing points  $z$  fully cover the space of the categories, implying that the learned inducing points meaningfully cover the distribution of the textual embeddings.

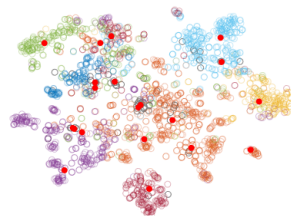


Figure 2: t-SNE visualization of learned network structural embeddings of Cora. Unfilled circles are individual nodes and inducing points are denoted as red dots.