# **Deep Generative Models for Subgraph Prediction**

Erfaneh Mahmoudzadeh <sup>a,\*</sup>, Parmis Naddaf <sup>a</sup>, Kiarash Zahirnia <sup>a</sup> and Oliver Schulte <sup>a</sup>

<sup>a</sup>School of Computing Science, Simon Fraser University

Abstract. Graph Neural Networks (GNNs) are important across different domains, such as social network analysis and recommendation systems, due to their ability to model complex relational data. This paper introduces subgraph queries as a new task for deep graph learning. Unlike traditional graph prediction tasks that focus on individual components like link prediction or node classification, subgraph queries jointly predict the components of a target subgraph based on evidence that is represented by an observed subgraph. For instance, a subgraph query can predict a set of target links and/or node labels. To answer subgraph queries, we utilize a probabilistic deep Graph Generative Model. Specifically, we inductively train a Variational Graph Auto-Encoder (VGAE) model, augmented to represent a joint distribution over links, node features and labels. Bayesian optimization is used to tune a weighting for the relative importance of links, node features and labels in a specific domain. We describe a deterministic and a sampling-based inference method for estimating subgraph probabilities from the VGAE generative graph distribution, without retraining, in zero-shot fashion. For evaluation, we apply the inference methods on a range of subgraph queries on six benchmark datasets. We find that inference from a model achieves superior predictive performance, surpassing independent prediction baselines with improvements in AUC scores ranging from 0.06 to 0.2 points, depending on the dataset.

## 1 Introduction: Subgraph Queries

Due to the importance and ubiquity of graph data, deep graph learning is a growing field with many applications. Typical graph learning tasks include link prediction (predicting the existence of a target link given as evidence other links and node attributes) and node classification (predicting a node label given the node links and labels of other nodes). This paper addresses subgraph prediction, which is a significant novel generalization of the link prediction and node classification tasks [9]. A subgraph query aims to estimate the probability of a target subgraph (i.e., a set of links, node labels, and node features), given an observed subgraph as evidence. Since subgraph components are not independent of each other, optimal subgraph prediction is different from predicting each target link/node label independently. We introduce VGAE+, an extension of the VGAE base generative model [13]. To answer subgraph queries, the trained VGAE+ models takes an attributed evidence graph as input and returns a probability for the target subgraph.

*Examples.* We provide some examples to illustrate the power of subgraph prediction, describe use cases, and clarify the relationship to link prediction and node classification. Figure 1 shows a scenario



Figure 1: Example of joint link prediction and node classification: predicting potential actors for Christopher Nolan's new movie and its genre. Target links are dashed lines marked with "?". Solid links are specified as evidence.

with a new film directed by Christopher Nolan. Previously, he collaborated with Cillian Murphy, Christian Bale, and Tom Hardy in "The Dark Knight Rises," and with Cillian Murphy in "Oppenheimer." Suppose our evidence stipulates that Cillian Murphy is cast in this new movie, and we want to predict which previous collaborators will join him in the movie, and which type of movie it will be. The subgraph prediction task is to jointly predict (i) the movie's genre and (ii) which additional actors will join the cast.

Predicting the genre of the new movie is an instance of *single node classification*, an extensively studied task [30, 2]. Predicting whether Tom Hardy will be in the new movie (independent of other actors) is an instance of *single link prediction*, an equally extensively studied task [9]. Predicting whether Tom Hardy, Cillian Murphy, and Christian Bale will all join the movie, is an instance of *joint link prediction*, a recently introduced task [16] that generalizes single link prediction generalizes both joint link prediction and node classification.

For another example, consider fraud detection in financial networks. A node classification approach is to label a network node as "suspicious" or "normal" [23]. However, recent attack types involve collusion among many network nodes controlled by the attackers [14]. Detecting collusion can be achieved by a more powerful approach based on joint node classification. Moreover, attack nodes exchange messages at a substantially higher frequency than normal nodes [27]. Subgraph prediction can capture joint association patterns among node labels and links to model both which nodes are suspicious and how suspicious nodes interact with each other.

**Motivation.** Subgraph queries are substantially more expressive than the previously studied link prediction and node classification tasks. They allow users to choose as prediction target a *set* of links

<sup>\*</sup> Corresponding Author. Email:ema61@sfu.ca.

and a *set* of node features or labels rather than assuming a fixed-size prediction target (e.g., single link, single label), instead of choosing one of them. Second, users can provide different kinds of evidence, depending on what is known, and use a single system for answering the query. In contrast, previous researchers have developed a separate customized method for different evidence types. Answering general relational queries, including subgraph queries, is a major use case for non-neural statistical-relational models, such as Markov Logic networks [3], which motivates our goal of expanding the prediction tasks to which neural models can be applied.

Recent generative AI models have shown the usefulness for many users of a single multi-task system. For deploying graph learning, a single query answering system is important in a production environment where we do not know in advance which graph queries will be important, and users may not have the resources to build a customized machine learning solution for different query types. In this work, we are proposing VGAE+, a generative model that supports inference to return a probabilistic answer to a subgraph query whose target and evidence set can be specified by the user. A generative model, unlike a discriminative model, supports prediction with targets of varying sizes.



**Figure 2**: After training a single GGM, the approximate inference methods described in this paper can answer a user query.

Approach. Figure 2 shows our system design. Our approach is a form of *domain-specific pre-training*: We learn a probabilistic Graph Generative Model (GGM) from data in a deployment domain. The model is used to answer queries with no further learning required. The deployment domain can be large (e.g., 1M nodes), whereas queries are typically small (e.g., involving 10-100 nodes). Specifically, we train a generative Variational Graph Auto-Encoder (VGAE) [13, 9] inductively, so that the VGAE can be applied to query targets of different sizes that may involve unseen nodes. The generative probabilities over graphs (implicitly) define a conditional probability P(target\_subgraph|evidence\_subgraph) for every subgraph query. The VGAE utilizes an encoder-decoder architecture with a Graph Neural Network (GNN) as an encoder, augmented to generate node labels/features as well as links. We describe and evaluate approximate inference methods for estimating conditional subgraph probabilities from a trained VGAE.

**Evaluation.** We evaluate a variety of query types on six benchmark homogeneous graphs, depending on whether 1) the target subgraph involves a single neighbor or multiple neighbors of a target node, and 2) the evidence comprises only links among training nodes (transductive), or also links that connect training and test nodes (inductive). Our main baseline is to estimate link probabilities and node labels independently using SOTA node classification/link prediction methods. For subgraph queries, we find that the VGAE model achieves superior or competitive performance on all datasets and settings, especially inductive inference. For node classification/link prediction problems, we find that inference from the augmented VGAE model is competitive with custom methods.

Contributions. Our main contributions are as follows.

- Introducing the task of conditional subgraph prediction.
- Approximate inference methods, including sampling, for answering subgraph queries from a single trained GGM.
- Introducing VGAE+, which is an augmented VGAE for jointly modeling links and node features/labels. VGAE+ is trained with a new objective that utilizes Bayesian optimization to weight link prediction and node classification.

Paper Organization. We review related work, then define subgraph queries. We describe inductive training of a VGAE based on the Evidence Lower BOund (*ELBO*) likelihood approximation and inference methods for answering subgraph queries from the trained VGAE. An extensive empirical evaluation compares subgraph prediction from a VGAE model to 6 baselines on 6 benchmark datasets.

#### 2 Related Works

Link Prediction. For deep graph models, the most common setting is single link completion: predict a single target link given a set of known links, which typically include a large set of training links. Other recent variants include the following single link tasks. (1) Condition on the attributes of the two target nodes only [10]. (2) Condition on links among other test nodes only [24]. (3) Condition on links among training and test nodes [35]. These single-link tasks are generalized by the recently introduced joint link prediction task [16]. In joint link prediction, the target is a set of links to be predicted jointly, given another set of links and node features as evidence. Joint link prediction is the closest predecessor to our work in that it aims to answer a large class of graph queries from a single model. However, joint link prediction does not provide the capability of predicting node labels or features, which is important in many applications. Subgraph prediction requires training a generative model to support the dual task of joint link prediction and joint node classification, which we address with Bayesian optimization in this paper.

*Node Classification.* Node classification is one of the most common tasks in graph analysis. The goal is to predict a class for each unlabelled node in the graph based on available graph evidence [12], which usually includes all links and node features. Graph neural networks (GNNs) achieve state-of-the-art performance [30]. We include a recent method [11] in our baseline comparisons. The common semi-supervised node classification task is similar to a subgraph query in that it requires a joint prediction of node labels. It is different in that a subgraph query requires predicting links as well.

*Two-Task GNNs vs. Inference from a Model.* There is a recent trend towards two-task GNNs that support both node classification and link prediction [29]. We include the recent two-task method GiGa-MAE among our baselines [21]. While two-task GNNs are designed to support both node classification and link prediction, they require separate training for each. In contrast, we train a single model for the large set of diverse tasks that can be represented as subgraph queries. To our knowledge, inference from a generative model to estimate subgraph probabilities is a new application of GGMs. The only other paper that addresses inference from a deep GGM is [16], which covers link prediction only. We develop a novel augmented variant of the VGAE to generate node labels and features as well as links. The

Graph Variational Auto Encoder (GVAE) [22] also generates node labels and features, but it cannot be applied to a single training graph. Also the GVAE does not utilize Bayesian optimization to balance the relative importance of modeling links vs. node labels/features.

As we mentioned in the introduction, statistical-relational models also support subgraph queries, but they are based on very different assumptions and model classes (e.g. exponential graph models) from deep GGMs, so we leave a direct comparison for future work.

*Graph Queries*. Powerful (non-probabilistic) graph query languages have been developed, such as Cypher and SPARQL [6, 19, 7]. Such graph queries typically return a set of nodes or links that satisfy a complex condition. In contrast, subgraph queries return a probability for a target subgraph.

Inductive Graph Training. Inductive graph learning has been a major topic in recent research [8, 33, 20]. To support link prediction for nodes not seen during training, our query answering approach can be used with any inductive encoder-decoder graph neural network architecture, trained with the variational *ELBO* objective described below. To our knowledge, this is the first use of the variational *ELBO* objective to support inductive subgraph prediction.

*Graph Generative Models.* Inference from a model requires the following properties of a GGM. (1) Supports the computation/estimation of explicit graph probabilities. (2) Admits a conditional variant. (3) Applies to graphs of different sizes. As we show in this paper, VGAEs meet these requirements, so we base our experiments on them. Specifically, we employ the most recent VGAE designed for link prediction [16]. Another advantage of the VGAE is that it is designed for a single large dataset, as are most methods for link prediction and node classification, so we can employ the same benchmark datasets. In contrast, other deep GGMs are usually trained on datasets with many disjoint graphs [4] (e.g., molecules). We believe that developing deep GGMs so that they support answering subgraph queries is a fruitful new direction for GNNs.

Subgraph Classification. Alsentzer et al. [1] identify subgraph prediction as an important task. Like our approach, their query targets and evidence are subgraphs. However, their system addresses only *classifying* the target subgraph with a single label. In contrast, we assign a subgraph probability.

*Structure-Conditioned Graph Generation* is a recently studied variant where the user specifies an induced subgraph (the evidence) and a conditional graph generation process completes it [26, Sec.D.1], [5]. Our definition of joint link prediction allows that neither evidence nor target links define a complete graph, and that they may share nodes, which raises special challenges (see Section 5.1). In addition, our aim is to output an explicit probability over a joint link assignment, not output a graph. The fact that this recent work has studied a related task beyond single component prediction supports our motivation for more complex queries.

### 3 Subgraph Queries

In this section we formally define our problem, answering probabilistic subgraph queries (SQs). SQs specify a set of links and attributes for a set of query nodes. The specification can involve a language as complex as first-order logic [3]. In this paper we introduce a relatively simple SQ syntax—essentially, conjunctions of links, and/or node attributes and labels—that suffices to express the prediction tasks that have been previously studied in graph learning (Section 2).

**Data Format.** An attributed labelled graph is a pair G = (V, E) comprising a finite set of nodes V with features and labels, and links E, which may be positive (present) or negative (absent). Each node is

assigned a k-dimensional attribute  $x_i$  with k > 0 and a node label  $l_i$ . A graph with N nodes can be represented by the following objects.

- An  $N \times N$  adjacency matrix A with  $\{0, 1\}$  Boolean entries.
- An  $N \times k$  node feature matrix **X** with  $\{0, 1\}$  Boolean entries.
- An N×l node label matrix L with a one-hot encoding of l discrete labels for each node.

A target subgraph of graph G = (V, E) is a graph  $G^Y = (V^Y, E^Y)$  where  $V^Y \subseteq V$  and  $E^Y \subseteq E$ . Similarly, an evidence subgraph of graph G is a graph  $G^E = (V^E, E^E)$  where  $E^E \subseteq E$  and  $E^E \cap E^Y = \emptyset$ . We refer to the nodes  $V^E$  that appear in the evidence as **evidence nodes**, to the nodes  $V^Y$  that appear in the target as **target nodes**, and to their union as **query nodes** (i.e.,  $V = V^E \cup V^Y$ ).

**Definition of Subgraph Queries.** A relational random variable corresponds to either a node attribute or an adjacency. Let  $\mathcal{A} = \{\mathbf{A}[u,v] : u \in V, v \in V\}$  be the set of **link variables**, and  $\mathcal{L} = \{\mathbf{L}[u] : u \in V\}$  be the set of **node label variables**, and  $\mathcal{X} = \{\mathbf{X}[u] : u \in V\}$  be the set of **feature variables**.

A subgraph query  $\mathbb{P}(target | evidence)$  is of the form

$$\mathbb{P}(\mathcal{A}^Y = a^Y, \mathcal{L}^Y = l^Y, \mathcal{X}^Y = x^Y | \mathcal{A}^E = a^E, \mathcal{L}^E = l^E, \mathcal{X}^E = x^E)$$
 where

- where
- $\mathcal{A}^{Y} = \{\mathbf{A}[u_{i}, v_{i}] : i = 1, ..., |\mathcal{A}^{Y}|, \mathbf{A}[u_{i}, v_{i}] \in E^{Y}\}$  is the list of binary target links, each assigned a value  $a_{i} \in \{0, 1\}$ . Similarly,  $\mathcal{A}^{E} = \{\mathbf{A}[u_{i}, v_{i}] : i = 1, ..., |\mathcal{A}^{E}|, \mathbf{A}[u_{i}, v_{i}] \in E^{E}\}$  is the list of binary evidence links.
- L<sup>Y</sup> = {L[u<sub>i</sub>] : i = 1,..., |L<sup>Y</sup>|, u<sub>i</sub> ∈ V<sup>Y</sup>} is the list of target node labels, each assigned a one-hot encoding of the *l* class labels l<sub>i</sub> ∈ {0,1}<sup>l</sup>. Similarly, L<sup>E</sup> = {L[u<sub>i</sub>] : i = 1,..., |L<sup>E</sup>|, u<sub>i</sub> ∈ V<sup>E</sup>} is the list of evidence node labels.
   X<sup>Y</sup> = {X[u<sub>i</sub>] : i = 1,..., |X<sup>Y</sup>|, u<sub>i</sub> ∈ V<sup>Y</sup>} is the list of
- $\mathcal{X}^{Y} = \{ X[u_i] : i = 1, ..., |\mathcal{X}^{Y}|, u_i \in V^{Y} \}$  is the list of target node features, each assigned a feature vector  $x_i \in \mathbb{R}^{1 \times k}$ . Similarly,  $\mathcal{X}^{E} = \{ X[u_i] : i = 1, ..., |\mathcal{X}^{E}|, u_i \in V^{E} \}$  is the list of evidence node features.

The indices *i* index elements in a query not nodes in a graph. Given a partition of nodes into observed training nodes and unobserved test nodes, a query is **inductive** if a test node appears in the target nodes.

Note that the target and evidence specifications can be and typically are *partial* in that some graph components are unspecified. For example, for two evidence nodes u and v, the evidence may specify their features, but not whether there exists a link between them or not. Our definition treats a node feature x[u] as a group, in that either all or no features of node u are specified in a query. This restriction is not essential; we use it mainly to simplify notation.

**Example** Figure 3(Right) shows an example of the most complex SQ type we consider in this research.

- The query nodes are 1–6.
- The query target subgraph comprises:
  - two positive target links defined by the pairs (1, 4), (5, 4).
  - two negative target links defined by the pairs (3, 4), (6, 4).
  - the labels of nodes 1, 3, 5, 6,
- The query evidence subgraph comprises:
  - five positive evidence links defined by the pairs (1,2), (1,3), (2,3), (3,6), (1,5).

- four negative evidence links defined by the pairs (3,5), (2,6), (2,5), (1,6).
- The features of all query nodes (not shown).
- The presence/absence of links (2, 4) and (6, 4) is unspecified.

The perfect query answer is to assign probability 1 to the ground truth subgraph shown on the Figure 3(Left). The query is inductive because nodes 4,5,6 are not observed during model training.



**Figure 3**: **Left**: Input graph with partition of nodes. Node labels are green and blue. **Right**: Neighborhood query: Target labels and links are colored red. Black dashed links are unspecified links in evidence.

#### 4 Variational Graph Auto-Encoder Training

We describe our generative model, a VGAE augmented with feature and label reconstruction, with training and implementation details. Figure 4 shows the VGAE training architecture.



Figure 4: Encoder-Decoder Training Architecture

#### 4.1 Augmented VGAE Generative Model

Let Z be an  $N \times d$  matrix that represents latent node embeddings. In the VGAE model, links are generated independently given node embeddings. Following the GraphVAE approach [22], we generate node classes and node features independently as well given node embeddings. We thus utilize three decoder models (see Figure 4):

$$p_{\theta}(\boldsymbol{A}|\boldsymbol{Z}) = \prod_{u,v} p_{\theta}(\boldsymbol{A}[u,v]|\boldsymbol{z}[u],\boldsymbol{z}[v])$$

$$p_{\psi}(\boldsymbol{X}|\boldsymbol{Z}) = \prod_{u} p_{\psi}(\boldsymbol{X}[u]|\boldsymbol{z}[u])$$

$$p_{\phi}(\boldsymbol{L}|\boldsymbol{Z}) = \prod_{u} p_{\phi}(\boldsymbol{L}[u]|\boldsymbol{z}[u])$$
(1)

where  $p_{\theta} : \mathbb{R}^d \times \mathbb{R}^d \to [0, 1]$  is a trainable link decoder, and  $p_{\psi}$  resp.  $p_{\phi}$  denotes a trainable feature decoder resp. label decoder.

The graph **encoder**  $q_{\phi}(Z|X, A)$  is implemented by a GNN that takes as input an attribute graph and returns latent node embeddings. For compatibility with baseline methods, the encoder does not receive node labels as input, but adding them is straightforward.

A VGAE+ is trained using the variational ELBO objective [13, 9]:

$$\mathcal{L}(\theta, \phi) = -E_{\mathbf{Z} \sim q_{\phi}(\mathbf{Z}|\mathbf{X}, \mathbf{A})} [\alpha \times \ln p_{\theta}(\mathbf{A}|\mathbf{Z}) + \beta \times \ln p_{\psi}(\mathbf{X}|\mathbf{Z}) + \gamma \times \ln p_{\phi}(\mathbf{L}|\mathbf{Z})] + KL(q_{\phi}(\mathbf{Z}|\mathbf{X}, \mathbf{A})||p(\mathbf{Z}))$$
(2)

where KL(.||.) is the Kullback-Leibler divergence between two probability distributions. Here  $\alpha$ ,  $\beta$  and  $\gamma$  are hyperparameters that weight the importance of different reconstruction tasks. These hyperparameters are found by Bayesian optimization [17]. The Optimizer tunes these hyperparameters so that they minimize the model reconstruction loss on the validation graph. For details please see the Appendix [15].

**Inductive Training** In order to answer inductive queries involving unseen nodes, we train the VGAE+ without node IDs. We randomly partition the nodes in the input graph into 1) 70% training nodes  $V_{tr}$ , 2) 20% inductive test nodes  $V_{te}$  and 3) 10% validation nodes. We train a VGAE+ model on the input subgraph induced by  $V_{tr}$ .

#### 4.2 Implementation

**Encoder** The graph **encoder**  $q_{\phi}(\boldsymbol{Z}|\boldsymbol{X}, \boldsymbol{A})$  is implemented by a GNN that takes an attributed labelled graph as input and returns node embeddings  $\boldsymbol{Z}$ . The node embeddings are independent, and for each node, represent a conditional Gaussian distribution, such that

$$q(\boldsymbol{z}[u]|\boldsymbol{X}, \boldsymbol{A}) \sim N(\boldsymbol{\mu}[u], \boldsymbol{\sigma}[u])$$
(3)

with mean and covariance  $\mu[u], \sigma[u]$  for node u.

**Link Decoder** As a strong link decoder, we utilize a Stochastic Block Model (SBM), which is defined as:

$$p_{\theta}(\boldsymbol{A}[u,v]|\boldsymbol{z}[u],\boldsymbol{z}[v]) = \boldsymbol{z}[u]^{\top} \Lambda \boldsymbol{z}[v]$$
(4)

where  $\Lambda \in \mathbb{R}^{d \times d}$  is the trainable *d*-block matrix in the SBM.

**Feature Decoder** Nodes features are reconstructed independently, given node embeddings. Our benchmark datasets comprise discrete features, so our *feature decoder* is of the form  $p_{\psi} : \mathbb{R}^d \to \{0, 1\}^k$ 

$$p_{\psi}(\boldsymbol{X}|\boldsymbol{Z}) = \prod_{u=1}^{N} \prod_{d=1}^{k} p_{\psi}(\boldsymbol{X}_{d}[u]|\boldsymbol{z}[u]) = \prod_{u=1}^{N} \prod_{d=1}^{k} \sigma(\tilde{x}_{d}[u]), \quad (5)$$

where  $\psi$  are the parameters of a fully connected neural net feature decoder that maps a node embedding  $\boldsymbol{z}[u]$  to a k-dimensional node feature reconstruction  $\tilde{x}[u]$ .

**Node Classifier** The node classifier  $p_{\phi} : \mathbb{R}^d \to \{0, 1\}^l$ , classifies the node labels independently under a one-hot encoding, given the same node embeddings:

$$p_{\phi}(\boldsymbol{L}|\boldsymbol{Z}) = \prod_{u=1}^{N} p_{\phi}(\boldsymbol{L}[u]|\boldsymbol{z}[u]) = \prod_{u=1}^{N} softmax(\tilde{l}[u])$$
(6)

where  $\phi$  are the parameters of a fully connected node classifier that maps a node embedding z[u] to a *l*-dimensional vector  $\tilde{l}[u]$ .

#### 5 Subgraph Inference from a VGAE Model

Consider a subgraph query Equation (1) with *n* query nodes  $\mathbb{P}(\mathcal{A}^{Y} = a^{Y}, \mathcal{L}^{Y} = l^{Y}, \mathcal{X}^{Y} = x^{Y} | E = e)$  where *E* is a list of evidence variables. We define two subgraph inference models, deterministic and Monte Carlo, for the VGAE+ generative model. Essentially, the deterministic method uses node embeddings computed deterministically from the evidence, whereas the MC method samples node embeddings conditional on the evidence.

#### 5.1 Inference Models

For a fixed set of n query node embeddings  $\boldsymbol{z}$ , the target probability is given by

$$\mathbb{P}_{\eta}(\mathcal{A}^{Y} = \boldsymbol{a}^{Y}, \mathcal{L}^{Y} = \boldsymbol{l}^{Y}, \mathcal{X}^{Y} = \boldsymbol{x}^{Y} | \boldsymbol{z})$$
(7)

and can be computed by multiplying the independent decoder probabilities Eqs. (4) to (6) parametrized by  $\eta$ . The **posterior distribution** over the *n* node embeddings is a conditional Gaussian distribution, such that  $p(\mathbf{z}[u]|\mathbf{E} = \mathbf{e}) \sim N(\mu[u], \sigma[u])$  with mean and covariance  $\mu[u], \sigma[u]$  for node *u*. Equation (10) below shows how to approximate the posterior using the GNN encoder. Let  $\mu_E$  be the mean node embeddings from the posterior distribution. **Deterministic inference** utilizes the posterior mean embeddings:

$$\mathbb{P}(\mathcal{A}^{Y}, \mathcal{L}^{Y}, \mathcal{X}^{Y} | \boldsymbol{E} = \boldsymbol{e}) \approx \mathbb{P}(\mathcal{A}^{Y}, \mathcal{L}^{Y}, \mathcal{X}^{Y} | \boldsymbol{\mu}_{\boldsymbol{E}}).$$
(8)

**Monte Carlo (MC) inference** samples *S* node embeddings from the posterior distribution, and averages the subgraph probabilities:

$$\mathbb{P}(\mathcal{A}^{Y}, \mathcal{L}^{Y}, \mathcal{X}^{Y} | \boldsymbol{E} = \boldsymbol{e}) \approx \frac{1}{S} \sum_{s=1}^{S} \mathbb{P}(\mathcal{A}^{Y}, \mathcal{L}^{Y}, \mathcal{X}^{Y} | \boldsymbol{z}^{s}) \qquad (9)$$

where  $\boldsymbol{Z}^{s} \sim \prod_{u=1}^{n} p(\boldsymbol{z}[u] | \boldsymbol{E} = \boldsymbol{e}).$ 

#### 5.2 Computing the Posterior Distribution

Training a VGAE+ model M provides decoder models  $p_{\eta}$  and an encoder  $q_{\phi}$ . The challenge is that the encoder assumes as input an adjacency matrix, or equivalently, a complete subgraph induced by the query nodes. We bridge the gap between a partially specified subgraph and a complete induced subgraph by imputing graph components missing from the evidence with 0 as a default value. For missing links, using the 0 default is the approach taken in previous work [16, 13]. As discussed by [16], 0 default is appropriate for links: First, a message-passing encoder treats 0 links as uninformative, which is appropriate for for unspecified links. Second, because of graph sparsity, the mode of the true link posterior given the evidence is close to 0. For node features and labels, the 0 default is also appropriate given our one-hot encoding, since the VGAE+ encoder is trained not to propagate information from 0-valued features/labels.

Formally, to obtain **evidence embeddings**, we approximate the posterior distribution using the trained encoder  $q_{\phi}$ :

$$p(\boldsymbol{Z}|\boldsymbol{E}=\boldsymbol{e}) \approx q_{\phi}(\boldsymbol{Z}|\boldsymbol{A}^{\boldsymbol{E},\boldsymbol{0}},\boldsymbol{X}^{\boldsymbol{E},\boldsymbol{0}})$$
(10)

where the evidence matrices  $A^{E,0}$  and  $X^{E,0}$  are defined as follows: 1) The dimension of  $A^{E,0}$  is  $n \times n$ . If  $\mathcal{A}^{E}$  specifies a link assignment  $\mathcal{A}[u, v] = e_i$ , then  $A^{E,0}[u, v] := e_i$ ; otherwise  $A^{E,0}[u, v] := 0$ . 2) The dimension of  $X^{E,0}$  is  $n \times k$ . If  $\mathcal{X}^{E}$  specifies a feature vector  $\mathcal{X}[u] = x_i$ , then  $X^{E,0}[u] := x_i$ ; otherwise  $X^{E,0}[u] := 0$ . Similarly we can assign a default value of 0 to unspecified node labels, but our experiments do not utilize queries that include node labels as evidence. Note that while the number of training nodes may be very large, the number of query nodes is typically small (on the order of 10-100). Applying the same encoder to subgraphs of different sizes is possible because we train the VGAE+ model inductively. (For details please see the Appendix [15].)

To illustrate, consider Figure 3. Since the link between nodes 1 and 2 is specified to exist, the evidence adjacency matrix assigns  $A^{E,0}[1,2] := 1$ . The link between nodes 3 and 5 is specified not

to exist, so the evidence matrix assigns  $A^{E,0}[3,5] := 0$ . Since the link between nodes 5 and 6 is unspecified, the evidence matrix assigns  $A^{E,0}[5,6] := 0$ . If the feature vector for node 2 is unspecified, the evidence matrix assigns the zero feature vector:  $X^{E,0}[2] := 0$ .

#### 6 Evaluation

We detail our methodology then discuss our empirical results. Our GitHub repository provides a PyTorch implementation and datasets.<sup>1</sup>

#### 6.1 Experimental Design.

We describe our benchmark datasets, the design of the test queries, and how evaluation metrics are computed.

**Datasets** We utilize datasets from previous studies of GGMs [13, 32, 10]. Cora, ACM, and CiteSeer are citation networks, IMDb is a movies dataset, and Photo and Computers are co-occurrence networks based on Amazon data. The appendix presents dataset statistics [15].

**Data Preprocessing** Following previous link prediction studies [13], we add self-loops and make all links undirected (i.e., if the training data contains an adjacency,  $v \rightarrow u$ , it also contains  $u \rightarrow v$ .) Cora, CiteSeer, Photo, and Computers are homogeneous datasets, whereas ACM and IMDb are heterogeneous datasets. Since our comparison methods use homogeneous GNNs, we homogenize different non-hierarchical edge types, such that every edge in the adjacency matrix is represented by 0 for no link and 1 for link existence.

**Test Query Design.** We explain next our method for generating test subgraph queries of different types. The appendix contains visual examples of the test queries [15]. In all our test queries, the evidence E does not contain node labels, but specifies 1) the node feature vector for each query node, 2) the non-target links from the input graph. To define query target nodes, we randomly select a set of 100 test nodes as target nodes from the test nodes  $V_{te}$  (cf. Section 4.1).

**Single Neighbor Query** For each target node u, we randomly select two test links, one positive pair  $(u, v_+)$  from the neighborhood of u, one negative pair  $(u, v_-)$  from outside the neighborhood. The resulting query is of the form  $P(L[u], A[u, v_+], A[u, v_-]|E)$ : Which of the two links is true and what is the label of the target node? In inductive learning, the paired nodes  $v_+, v_-$  are both test nodes. In semi-inductive learning, each paired node  $v_+, v_-$  may be either a training node or a test node.

**Neighborhood Query** For each target node u with at least one neighbor, let  $v_1^+, \ldots, v_{deg(u)}^+$  enumerate the nodes in the neighborhood of u. We randomly select negative test nodes  $v_1^-, \ldots, v_{deg(u)}^-$ . The resulting query is of the form  $\mathbb{P}(\{\boldsymbol{L}[v_i^+], \boldsymbol{A}[u, v_i^+], \boldsymbol{L}[v_i^-], \boldsymbol{A}[u, v_i^-] : i = 1, \ldots, deg(u)\}|\boldsymbol{E}) :$ which nodes are neighbors and what are their labels? (cf. Figure 1.) In semi-inductive learning, the paired nodes  $v_1^-$  are training nodes. In inductive learning, they can be either test nodes or training nodes.

**Evaluation Metrics.** For single-neighbor queries, we compute the mean of each metric across all test queries. For neighborhood queries, the metric is calculated separately for each query, over all the components of the query (e.g., all the neighbors to be predicted). We then report the mean of each metric across all neighborhood queries. We separately score link prediction and node classification using ROC-AUC. ROC-AUC is a good metric for evaluating both

<sup>&</sup>lt;sup>1</sup> https://github.com/erfmah/Answering\_Graph\_Queries

link prediction and node classification because it provides a comprehensive and threshold-independent measure of model performance, making it more robust to class imbalance than the accuracy metric. Separate scoring is supported by the VGAE+ inference model (7), but in fact favors the single-component prediction baselines over our joint prediction model. The reason for independent scoring is that our baseline methods are not designed for joint prediction, so we use them only to obtain scores for individual components. The joint predictive performance for a subgraph query is measured by the average AUC score over predicted links and node labels. Other metrics (Hit-Rate@20% and F1-macro) are in the appendix [15].

#### 6.2 Query Answering Comparison Methods.

**Our VGAE+ approach** We built on the VGAE implementation for joint link prediction [16], following the authors' recommendation of using MC inference for link prediction and a graph isomorphism network (GIN) encoder [31] with 2 layers and embedding dimension 128. We extended their VGAE architecture to include the reconstruction of node features and labels, as described in Section 4. Node labels are predicted using the node classifier Eq. (6), trained end-to-end. We give results for deterministic prediction Eq. (8) and MC prediction Eq. (9) with 30 samples.

We select baselines that are well established and represent a variety of approaches to node classification or single link prediction. Node classification methods perform joint label prediction (known as semi-supervised node classification [30]) like our VGAE+ model. We apply single-link prediction methods to a set of target links separately to obtain a probability for each target link (cf. [16]). We organize methods according to whether they were originally evaluated on both link prediction and node classification, or just one or the other.

**Two-Task Methods** train two separate link prediction and node classification models using the same GNN architecture.

**Generalizable Graph Masked AutoEncoder (GiGaMAE)** [21] introduces a novel graph encoder based on aligning different node embeddings that respectively encode structural and feature information. We use the authors' code to obtain node label and link prediction probabilities. We train GiGaMAE inductively for inductive queries.

**GraphSage** [8] is an inductive model. We have used the DGL implementation [28], with the supervised training mode, for node classification, where the decoder sees the node labels. For link prediction, we train an SBM link decoder end-to-end using their unsupervised training. **Graph Attention Networks (GAT)** are a kind of GNN that computes different weights for different nodes in a neighborhood [25]. We used the DGL implementation (GAT) [28], in supervised training mode, for node classification, where the decoder is exposed to the node labels. The GAT embedding system was trained with the same micro- $F_1$  score as in the original paper [25]. For link prediction, we train an SBM link decoder end-to-end using unsupervised training.

**Link prediction baselines** include the following. **SEAL** is a wellknown method for transductive link prediction [34]. We train SEAL inductively by omitting node ids (cf. [24, 16]). The basic idea of SEAL is to embed the subgraph around a target link. **DEAL** was designed for both transductive and inductive cold-start link prediction tasks [10]. The distinguishing feature of the DEAL approach is that at test time, it uses only the node attributes for link prediction. This makes it a strong baseline for independent prediction since it is not sensitive to the presence of evidence links. **Node classification baselines** utilize the following encoders. **MV-GRL** is an inductive self-supervised approach for learning representations of nodes and graphs by contrasting different structural views of graphs [11]. We use the authors' code to train a node classification model. **DeepWalk** is a transductive graph embedding method that learns node representations by treating random walks on the graph as sentences and applying word embedding techniques [18]. It captures structural information of the graph by mapping nodes to low-dimensional vectors in a continuous space. The reported results are based on the DGL implementation [28]. Since DeepWalk cannot be adapted for inductive queries, we trained it transductively on the complete input graph. Node classification is performed by a logistic regression model [18].

#### 6.3 Experimental Results

We discuss our results for different types of test queries, starting with the most complex subgraph queries. It is important to note that *while VGAE+ addresses link prediction and node classification simultaneously, the baseline models are customized for each task separately.* The VGAE+ is trained once for all subgraph queries, whereas the baseline models are trained for a specific task. Because the VGAE+ model is solving a harder problem, matching the performance of the baseline methods is a good result. Most baselines were developed for transductive graph prediction, some only transductively, some both transductively and inductively. We therefore expect them to perform relatively better in the semi-inductive than in the (fully) inductive setting. As we will see, our results confirm this expectation.

#### 6.3.1 Subgraph Queries.

The reported AUC results for these baseline models are the average AUC scores for both link prediction and node classification tasks. We do not show standard deviations for simplicity, they range from 0.1 to 0.2. Table 1a shows the results for subgraph queries, which predict *both* links and node labels. In the inductive setting, all target nodes, and potentially some of their neighbors, are unseen during training.

**Neighborhood Queries** Considering the *inductive* setting, *both VGAE+ methods score more highly than both two-task baselines on almost all datasets.* The exceptions are the ACM and IMDb datasets, where the VGAE-Det score is similar to the GiGaMAE baseline, whereas VGAE-MC is clearly the best predictor. The biggest improvement for VGAE-MC is observed on the IMDb dataset (0.16 AUC points). The strong performance on inductive queries illustrates how a generative model can generalize to data that may be incomplete or new, by inferring unseen information.

In the *semi-inductive* setting, where some target nodes are observed during training, VGAE-MC achieves the highest score on 4 out of 6 datasets, and a competitive score on IMDb. The VGAE-MC top scores are substantially higher than the baselines; for instance on Computers it outperforms by more than 0.2 AUC points.

Our results show that for *some datasets, inference through sampling from an approximate posterior yields substantive accuracy improvements.* However, sampling takes longer to produce query results. To give a sense of the difference, for the biggest *Computers* dataset, sampling 30 node embeddings vs. using a single deterministic one takes about 4 times as long (39.02s/query for VGAE-MC vs. 11.951s/query for VGAE-Det), using a single NVIDIA A40 GPU.

#### Table 1: AUC scores on Test Queries

(a) AUC results for subgraph queries in semi-inductive and inductive settings.

-		Single Neighbor Queries							Neighborhood Queries						
		Cora	ACM	IMDb	CiteSeer	Photo	Computers		Cora	ACM	IMDb	CiteSeer	Photo	Computers	
Semi-inductive	VGAE-Det VGAE-MC GAT	$\frac{0.937}{0.949}$	$-\frac{0.801}{0.828}$	$-\frac{0.863}{0.900}$	$\begin{array}{r} 0.949 \\ - \ \underline{0.941} \\ 0 \ \overline{910} \end{array}$	<b>0.960</b> 0.951 0.935	0.910 $\frac{0.960}{0.922}$		0.868 <b>0.933</b> 0.829	0.700 0.716 0.726	0.764 - $0.734$ - $0.733$	$-\frac{0.934}{0.951}$	0.854 <b>0.949</b> -0.881	$\frac{0.789}{0.918} \frac{0.789}{0.840}$	
	GraphSage GiGaMAE	0.797 0.913	0.723 0.883	0.748 <b>0.920</b>	0.750 0.896	0.860 <u>0.956</u>	0.858 0.948		0.701 0.917	0.689 0.782	0.726 <u>0.746</u>	0.786 0.896	0.846 <u>0.900</u>	$\frac{0.010}{0.724}$ 0.703	
Inductive	VGAE-Det VGAE-MC GAT GraphSage GiGaMAE	<b>0.932</b> - <u>0.925</u> <u>0.728</u> 0.810 0.896	$\begin{array}{r} 0.751 \\ - \underbrace{0.832}_{\overline{0.841}} \\ - \underbrace{0.783}_{\overline{0.668}} \end{array}$	0.880 0.900 0.760 0.644 0.715	<b>0.963</b> <u>0.912</u> <u>0.798</u> 0.702 0.896	0.947 0.953 0.849 0.761 0.788	$ \frac{\frac{0.945}{0.954}}{\frac{0.866}{0.712}}$		0.825 <b>0.866</b> 0.681 0.710 <u>0.828</u>	0.697 <b>0.781</b> 0.573 0.475 <u>0.716</u>	$-\frac{\frac{0.823}{0.865}}{\frac{0.650}{0.463}}$	<b>0.943</b> <u>0.889</u> <u>0.648</u> 0.709 0.885	<b>0.884</b> <u>0.825</u> <u>0.630</u> 0.684 0.816	<b>0.885</b> <u>0.819</u> <u>0.582</u> 0.566 0.718	

(b) AUC results for link prediction in semi-inductive and inductive settings.

		Single Link Queries							Joint Link Queries						
		Cora	ACM	IMDb	CiteSeer	Photo	Computers		Cora	ACM	IMDb	CiteSeer	Photo	Computers	
	VGAE-Det	0.898	<u>0.956</u>	0.937	0.941	0.947	0.871		0.759	<u>0.962</u>	0.894	0.910	0.798	0.753	
Comi industivo	VGAE-MC	0.912	0.954	0.957	0.952	0.956	0.935		0.886	0.960	0.901	0.941	0.940	0.928	
Semi-inductive	SEAL	0.910	0.904	-0.906-	0.948 -	0.935	0.908		0.741	0.500	0.620	- 0.752	0.684	0.701	
	GAT	0.862	0.943	0.874	0.913	0.902	0.892		0.764	0.925	0.845	0.904	0.872	0.814	
	GraphSage	0.854	0.969	0.917	0.589	0.858	0.913		0.725	0.962	0.893	0.666	0.853	0.823	
	GiGaMAE	0.907	0.943	0.950	0.950	0.949	0.917		0.902	0.959	0.963	0.940	0.930	0.906	
	DEAL	0.800	0.986	0.981	0.914	0.832	0.823		0.676	0.979	0.962	0.910	0.857	0.826	
	VGAE-Det	0.891	0.963	0.939	0.950	0.936	0.925		0.690	0.954	0.867	0.925	0.857	0.821	
	VGAE-MC	0.865	0.970	0.942	0.900	0.931	0.922		0.778	0.968	0.935	0.856	0.760	0.752	
Inductive	SEAL	0.756	-0.666	0.854	0.627 -	0.924	0.972		0.693	0.681	0.830	- 0.679	0.506	0.450	
	GAT	0.667	0.870	0.796	0.731	0.796	0.835		0.491	0.699	0.632	0.470	0.453	0.380	
	GraphSage	0.561	0.593	0.492	0.504	0.552	0.444		0.543	0.535	0.467	0.512	0.528	0.507	
	GiGaMAE	0.887	0.732	0.929	0.943	0.798	0.764		0.725	0.920	0.918	0.921	0.854	0.706	
	DEAL	0.780	<u>0.902</u>	0.957	0.861	0.852	0.844		0.678	0.953	0.935	0.837	0.759	0.757	

(c) AUC results for node classification in semi-inductive and inductive settings.

				Single	Node Quer	ies		Joint Node Queries					
		Cora	ACM	IMDb	CiteSeer	Photo	Computers	Cora	ACM	IMDb	CiteSeer	Photo	Computers
	VGAE-Det	0.977	0.615	0.789	0.957	0.973	0.948	0.977	0.425	0.634	0.959	0.910	0.824
Comi industivo	VGAE-MC	0.986	0.701	0.842	0.930	0.946	0.986	0.980	0.471	0.568	0.960	0.959	0.909
Semi-inductive	GAT	0.861	$-0.876^{-}$	0.813	0.907	0.967		0.893	- 0.526	0.621	- 0.896	0.889	0.866
	GraphSage	0.740	0.477	0.580	0.911	0.863	0.804	0.677	0.416	0.460	0.907	0.840	0.626
	GiGaMAE	0.920	0.823	0.890	0.842	<u>0.963</u>	0.941	0.932	<u>0.604</u>	0.529	0.852	0.871	0.500
	MVGRL	0.888	0.708	0.788	0.807	0.963	<u>0.980</u>	0.853	0.715	0.767	0.815	<u>0.953</u>	0.892
	Deep Walk	0.868	0.643	0.684	0.847	0.950	0.862	0.726	0.567	0.553	0.731	0.937	0.819
	VGAE-Det	0.974	0.540	0.821	0.976	0.958	0.965	0.961	0.441	0.780	0.961	0.912	0.890
	VGAE-MC	0.986	0.693	0.856	0.924	0.976	0.987	0.954	0.595	0.796	0.922	0.890	<u>0.886</u>
Inductive	GAT	0.790	0.812	0.724	0.865	0.902	0.896	0.870	0.446	0.668	0.826	0.806	0.784
	GraphSage	0.970	0.973	0.797	0.900	<u>0.970</u>	<u>0.980</u>	0.877	0.416	0.460	0.907	0.840	0.626
	GiGaMAE	0.906	0.604	0.501	0.850	0.778	0.847	0.932	0.513	0.482	0.850	0.778	0.730
	MVGRL	0.650	0.804	0.684	0.550	0.835	0.801	0.614	0.679	0.632	0.534	0.892	0.827

**Single Neighbor Queries** The task is to predict the link and node label for a single neighbor of a target node. VGAE-MC scores higher than both baselines *in the inductive setting*. The improvement over the next best baseline is substantive on 3 out of 6 datasets; for instance on IMDb, the improvement is almost 0.2 AUC points.

GiGaMAE was designed for the transductive learning, and is generally competitive with VGAE-MC on single neighbor semiinductive queries. VGAE-MC scores substantially higher than GraphSage, for instance on CiteSeer by almost 0.2 AUC points.

In conclusion, inference from a VGAE model offers the best performance for predicting jointly both links and node labels. We next break down the subgraph AUC scores into their link prediction and node label components. This allows us to understand the results in more detail and to compare with more single-task baselines.

#### 6.3.2 Link Prediction Queries.

Table 1b shows the AUC scores for link prediction queries. For *inductive* joint link prediction, VGAE-MC achieves top scores on 4 out 6 datasets compared to the baselines. GiGaMAE scores better on Photos and CiteSeer, although not as high as VGAE-Det inference. VGAE-MC outperforms the link prediction baselines other

than DEAL and GiGaMAE. For *semi-inductive* joint link prediction, DEAL and GiGaMAE outperform VGAE-MC on 3 out of 6 datasets, but only by 0.02 AUC points at most (on Cora).

For *single-link inductive queries*, VGAE-MC outperforms the strongest baselines GiGaMAE and DEAL on 3 out of 6 datasets (ACM, Photo and Computers) by 0.08 AUC points or more. On the remaining 3 datasets, the VGAE-MC score is competitive. VGAE+ inference beats the other single-link prediction baselines, except for Photo and Computers, where SEAL performs exceptionally well. For *single-link semi-inductive queries* the VGAE-MC scores are the highest on 4 out of 6 datasets, but similar in magnitude to the GiGa-MAE and DEAL baselines.

In conclusion, inference from a VGAE model offers excellent link prediction accuracy across a range of query tasks, including inductive and joint link prediction.

#### 6.3.3 Node Classification Queries.

Table 1c shows AUC scores for node classification queries. For *inductive joint* node classification, VGAE-MC achieves a higher score than the baseline methods on 4 out of 6 datasets, especially on IMDb (at least 0.16 AUC points) and Cora (at least 0.07 AUC points). On ACM the strong node classification baseline MVGRL achieves a higher score. Similarly in the *semi-inductive joint* node classification setting, VGAE+ has a higher score than the baselines on 4 out of 6 datasets. The improvement is strongest on CiteSeer (0.06 over GraphSage). MVGRL is exceptionally strong in the semi-inductive setting on ACM and IMDb.

For *inductive single* node classification, VGAE-MC again achieves the top score on 4 out of 6 datasets. GraphSage was designed for inductive node classification and is accordingly a strong baseline, especially on the ACM dataset. However, VGAE-MC outscores GraphSage by 0.06 AUC points on the IMDb dataset.

For *semi-inductive single* node classification, VGAE-MC achieves the top score on 3 out of 6 datasets. The biggest improvement is observed on Cora, 0.06 AUC points over GiGaMAE.

In conclusion, inference from a VGAE model offers strong node classification accuracy across a range of settings, especially for joint node classification.

Subgraph Queries vs. Node Classification and Link Prediction. Reviewing the results of inductive subgraph prediction in terms of link prediction and node classification, we observe that the high score of VGAE-MC on IMDb is due to its excellent score on *joint node classification*. On the ACM dataset, VGAE-MC achieves strong subgraph prediction through a high *joint link prediction score*. The strong AUC score on Computers is due to both high joint node classification scores and high joint link prediction scores.

Overall our experiments provide strong evidence that inference from an augmented VGAE model, *achieves an excellent balance between predicting links and node labels across different query types.* For single query types (e.g., link prediction), predictive performance is very competitive with custom baselines.

Ablation Study. Table 2 examines the importance of the components of our new training objective Eq. (2) for the augmented VGAE. The second row ( $\beta = 0$ ) shows that not reconstructing the node features leads to worse scores, especially on the IMDb and Computers datasets. This is remarkable since the neighborhood queries do not contain node features as a target. Turning off label reconstruction ( $\gamma = 0$ ) or link reconstruction ( $\alpha = 0$ ) leads to bad scores.

 Table 2: Ablation Study on the training objective Eq. (2).

	Inductive Neighborhood Query											
	Cora	ACM	IMDb	CiteSeer	Photo	Computers						
VGAE-Det	0.825	0.697	0.823	0.943	0.884	0.885						
$(\beta = 0)$	0.810	0.616	0.660	0.927	0.865	0.748						
$(\beta = 0, \gamma = 0)$	0.611	0.652	0.710	0.688	0.653	0.574						
$(\beta=0,\alpha=0)$	0.656	0.417	0.483	0.686	0.667	0.580						

#### 7 Conclusion, Limitations, and Future Work

A subgraph prediction query asks for the probability of a target subgraph, given the information from an evidence subgraph. Supporting inference to answer subgraph queries (SQs) is a new use case for a deep Graph Generative Model (GGM). Such a query answering system facilitates applying graph prediction in a production environment where multiple users pose a range of queries to be answered. In this paper we showed how inference from a trained Variational Graph Auto-Encoder (VGAE) model, augmented with feature/label decoders, can be used to answer SQs, in zero-shot manner without retraining the model. Bayesian optimization was effective in balancing the relative importance of modeling links and node features/labels in a dataset-dependent manner. We carried out an empirical evaluation on six benchmark datasets and a range of test queries. The application of joint prediction from a single VGAE yielded higher accuracy than baseline methods that predict graph components independently. The strong performance of VGAE+ highlights the value of using both node features and graph structure as co-training objectives that improve both node classification and link prediction.

*Limitations*. A limitation of our evaluation is that we considered only homogeneous graphs with a single link type. Deterministic and MC inference can be extended straightforwardly to knowledge graphs using a relational VGAE model.

*Future Work.* While the VGAE is a well-established GGM for a single training graph, other GGMs, especially auto-regressive and diffusion models, are known to have greater modeling power to capture complex correlation patterns in graphs [16]. Leveraging the greater expressive power of these GGMs to improve subgraph predictions over our strong VGAE baseline is a fruitful direction for future research, especially if they can be trained on single graph inputs. Extending subgraph prediction to more complex graphs, such as weighted and/or dynamic graphs, is a fruitful topic for future work. Another valuable direction is to apply inference from a model to find the most likely subgraph given evidence.

#### References

E. Mahmoudzadeh et al. / Deep Generative Models for Subgraph Prediction

3136

- [31] K. Xu, W. Hu, J. Leskovec, and S. Jegelka. How powerful are graph neural networks? ICLR, 2019.
- [1] E. Alsentzer, S. Finlayson, M. Li, and M. Zitnik. Subgraph neural networks. NeurIPS, 2020.
- [2] H. Cai, V. W. Zheng, and K. C.-C. Chang. A comprehensive survey of graph embedding: Problems, techniques, and applications. IEEE transactions on knowledge and data engineering, 2018.
- [3] P. Domingos and D. Lowd. Markov Logic: An Interface Layer for Artificial Intelligence. Morgan and Claypool Publishers, 2009.
- [4] F. Faez, Y. Ommi, M. S. Baghshah, and H. R. Rabiee. Deep graph generators: A survey. IEEE Access, 2021.
- F. Faez, N. Hashemi Dijujin, M. Soleymani Baghshah, and H. R. Ra-[5] biee. Scgg: A deep structure-conditioned graph generative model. Plos one, 2022.
- [6] N. Francis, A. Green, P. Guagliardo, L. Libkin, T. Lindaaker, V. Marsault, S. Plantikow, M. Rydberg, P. Selmer, and A. Taylor. Cypher: An evolving query language for property graphs. ACM, 2018.
- M. Galkin, Z. Zhu, H. Ren, and J. Tang. Inductive logical query answer-[7] ing in knowledge graphs. NeurIPS, 2022.
- W. Hamilton, Z. Ying, and J. Leskovec. Inductive representation learning on large graphs. NeurIPS, 2017.
- [9] W. L. Hamilton. Graph representation learning. Morgan & Claypool Publishers, 2020.
- Y. Hao, X. Cao, Y. Fang, X. Xie, and S. Wang. Inductive link prediction [10] for nodes having only attribute information. IJCAI, 2020.
- [11] K. Hassani and A. H. Khasahmadi. Contrastive multi-view representation learning on graphs. ICML, 2020.
- [12] P. Kazienko and T. Kajdanowicz. Label-dependent node classification in the network. Neurocomputing, 2012.
- [13] T. N. Kipf and M. Welling. Variational graph auto-encoders. NeurIPS, 2016.
- [14] Z. Li, A. L. G. Rios, and L. Trajković. Machine learning for detecting anomalies and intrusions in communication networks. IEEE Journal on Selected Areas in Communications, 2021.
- [15] E. Mahmoudzadeh, P. Naddaf, K. Zahirnia, and O. Schulte. Deep generative models for subgraph prediction, 2024. URL https://arxiv.org/abs/ 2408.04053
- [16] P. Naddaf, E. Mahmoudzaheh Ahmadi Nejad, K. Zahirnia, M. Jaeger, and O. Schulte. Joint link prediction via inference from a model. CIKM, 2023.
- Bayesian Optimization: Open source constrained [17] F. Nogueira. global optimization tool for Python, 2014-. URL https://github.com/ bayesian-optimization/BayesianOptimization.
- [18] B. Perozzi, R. Al-Rfou, and S. Skiena. Deepwalk: Online learning of social representations. ACM, 2014.
- [19] E. Prud'hommeaux and A. Seaborne. SPARQL Query Language for RDF. W3C Recommendation, 2008. URL http://www.w3.org/TR/ rdf-sparql-query/.
- [20] R. A. Rossi, R. Zhou, and N. K. Ahmed. Deep inductive graph representation learning. IEEE, 2018.
- [21] Y. Shi, Y. Dong, Q. Tan, J. Li, and N. Liu. Gigamae: Generalizable graph masked autoencoder via collaborative latent space reconstruction. ACM, 2023.
- [22] M. Simonovsky and N. Komodakis. Graphvae: Towards generation of small graphs using variational autoencoders. ICANN, 2018.
- [23] J. Tang, F. Hua, Z. Gao, P. Zhao, and J. Li. Gadbench: Revisiting and benchmarking supervised graph anomaly detection. NeurIPS, 2024.
- [24] K. Teru, E. Denis, and W. Hamilton. Inductive relation prediction by subgraph reasoning. ICML, 2020.
- [25] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio. Graph attention networks. ICLR, 2018.
- C. Vignac, I. Krawczuk, A. Siraudin, B. Wang, V. Cevher, and [26] P. Frossard. Digress: Discrete denoising diffusion for graph generation. In Proceedings of the Eleventh International Conference on Learning Representations, 2023.
- [27] J. Vitorino, I. Praça, and E. Maia. Sok: Realistic adversarial attacks and defenses for intelligent network intrusion detection. Computers & Security, 2023
- [28] M. Wang, D. Zheng, Z. Ye, Q. Gan, M. Li, X. Song, J. Zhou, C. Ma, L. Yu, Y. Gai, T. Xiao, T. He, G. Karypis, J. Li, and Z. Zhang. Deep graph library: A graph-centric, highly-performant package for graph neural networks. arXiv preprint arXiv:1909.01315, 2019.
- [29] Z. Wu, M. Zhan, H. Zhang, Q. Luo, and K. Tang. Mtgcn: A multitask approach for node classification and link prediction in graph data. Information Processing & Management, 2022.
- [30] S. Xiao, S. Wang, Y. Dai, and W. Guo. Graph neural networks in node classification: survey and evaluation. Machine Vision and Applications, 2022.

- [32] S. Yun, M. Jeong, R. Kim, J. Kang, and H. J. Kim. Graph transformer networks. NeurIPS, 2019.
- [33] H. Zeng, H. Zhou, A. Srivastava, R. Kannan, and V. Prasanna. Graph-SAINT: Graph sampling based inductive learning method. ICLR, 2020.
- [34] M. Zhang and Y. Chen. Link prediction based on graph neural networks. NeurIPS, 2018.
- [35] M. Zhang and Y. Chen. Inductive matrix completion based on graph neural networks. arXiv preprint arXiv:1904.12058, 2019.