

Representation Learning for Dynamic Graphs: A Survey

Seyed Mehran Kazemi

MEHRAN.KAZEMI@BOREALISAI.COM

Rishab Goel

RISHAB.GOEL@BOREALISAI.COM

Borealis AI, 310-6666 Saint Urbain, Montreal, QC, Canada

Kshitij Jain

KSHITIJ.JAIN@BOREALISAI.COM

Ivan Kobyzev

IVAN.KOBYZEV@BOREALISAI.COM

Akshay Sethi

AKSHAY.SETHI@BOREALISAI.COM

Peter Forsyth

PETER.FORSYTH@BOREALISAI.COM

Pascal Poupart

PASCAL.POUPART@BOREALISAI.COM

Borealis AI, 301-420 West Graham Way, Waterloo, ON, Canada

Editor: Karsten Borgwardt

Abstract

Graphs arise naturally in many real-world applications including social networks, recommender systems, ontologies, biology, and computational finance. Traditionally, machine learning models for graphs have been mostly designed for static graphs. However, many applications involve evolving graphs. This introduces important challenges for learning and inference since nodes, attributes, and edges change over time. In this survey, we review the recent advances in representation learning for dynamic graphs, including dynamic knowledge graphs. We describe existing models from an encoder-decoder perspective, categorize these encoders and decoders based on the techniques they employ, and analyze the approaches in each category. We also review several prominent applications and widely used datasets and highlight directions for future research.

Keywords: graph representation learning, dynamic graphs, knowledge graph embedding, heterogeneous information networks

1. Introduction

In the era of big data, a challenge is to leverage data as effectively as possible to extract patterns, make predictions, and more generally unlock value. In many situations, the data does not consist only of vectors of features, but also relations that form graphs among entities. Graphs arise naturally in social networks (users with friendship relations, emails, text messages), recommender systems (users and products with transactions and rating relations), ontologies (concepts with relations), computational biology (protein-protein interactions), computational finance (web of companies with competitor, customer, subsidiary relations, supply chain graph, graph of customer-merchant transactions), etc. While it is often possible to ignore relations and use traditional machine learning techniques based on vectors of features, relations provide additional valuable information that permits inference among nodes. Hence, graph-based techniques have emerged as leading approaches in the industry for application domains with relational information.

Traditionally, research has been done mostly on static graphs where nodes and edges are fixed and do not change over time. Many applications, however, involve dynamic graphs. For instance, in social media, communication events such as emails and text messages are streaming while friendship relations evolve. In recommender systems, new products, new users and new ratings appear every day. In computational finance, transactions are streaming and supply chain relations are continuously evolving. As a result, the last few years have seen a surge of works on dynamic graphs. This survey focuses precisely on dynamic graphs. Note that there are already many good surveys on static graphs (see, e.g., Hamilton et al. (2017b); Zhang et al. (2018b); Cai et al. (2018); Cui et al. (2018); Nickel et al. (2016a); Shi et al. (2016); Wang et al. (2017a)). There are also several surveys on techniques for dynamic graphs (see, e.g., Bilgin and Yener (2006); Zhang (2010); Spiliopoulou (2011); Aggarwal and Subbian (2014); Al Hasan and Zaki (2011)), but they do not review recent advances in neural representation learning.

We present a survey that focuses on recent representation learning techniques for dynamic graphs. More precisely, we focus on reviewing techniques that either produce time-dependent embeddings that capture the essence of the nodes and edges of evolving graphs or use embeddings to answer various questions such as node classification, event prediction/interpolation, and link prediction. Accordingly, we use an encoder-decoder framework to categorize and analyze techniques that encode various aspects of graphs into embeddings and other techniques that decode embeddings into predictions. We survey techniques that deal with discrete-and/or continuous-time events.

The survey is structured as follows. Section 2 introduces the notation and provides some background about static/dynamic graphs, inference tasks, and learning techniques. Section 3 provides an overview of representation learning techniques for static graphs. This section is not meant to be a survey, but rather to introduce important concepts that will be extended for dynamic graphs. Section 4 describes encoding techniques that aggregate temporal observations and static features, use time as a regularizer, perform decompositions, traverse dynamic networks with random walks, and model observation sequences with various types of processes (e.g., recurrent neural networks). Section 5 categorizes decoders for dynamic graphs into time-predicting and time-conditioned decoders and surveys the decoders in each category. Section 6 describes briefly other lines of work that do not conform to the encoder-decoder framework such as statistical relational learning, and topics related to dynamic (knowledge) graphs such as spatiotemporal graphs and the construction of dynamic knowledge graphs from text. Section 7 reviews important applications of dynamic graphs with representative tasks. A list of static and temporal datasets is also provided with a summary of their properties. Section 8 concludes the survey with a discussion of several open problems and possible research directions.

2. Background and Notation

In this section, we define our notation and provide the necessary background for readers to follow the rest of the survey. A summary of the main notation and abbreviations can be found in Table 1.

We use lower-case letters to denote scalars, bold lower-case letters to denote vectors, and bold upper-case letters to denote matrices. For a vector \mathbf{z} , we represent the i^{th} element of

the vector as $\mathbf{z}[i]$. For a matrix \mathbf{A} , we represent the i^{th} row of \mathbf{A} as $\mathbf{A}[i]$, and the element at the i^{th} row and j^{th} column as $\mathbf{A}[i][j]$. $\|\mathbf{z}\|_i$ represents norm i of a vector \mathbf{z} and $\|\mathbf{Z}\|_F$ represents the Frobenius norm of a matrix \mathbf{Z} . For two vectors $\mathbf{z}_1 \in \mathbb{R}^{d_1}$ and $\mathbf{z}_2 \in \mathbb{R}^{d_2}$, we use $[\mathbf{z}_1; \mathbf{z}_2] \in \mathbb{R}^{d_1+d_2}$ to represent the concatenation of the two vectors. When $d_1 = d_2 = d$, we use $[\mathbf{z}_1 \ \mathbf{z}_2] \in \mathbb{R}^{d \times 2}$ to represent a $d \times 2$ matrix whose two columns correspond to \mathbf{z}_1 and \mathbf{z}_2 respectively. We use \odot to represent element-wise (Hadamard) multiplication. We represent by \mathbf{I}_d the identity matrix of size $d \times d$. $\text{vec}(\mathbf{A})$ vectorizes $\mathbf{A} \in \mathbb{R}^{d_1 \times d_2}$ into a vector of size $d_1 d_2$. $\text{diag}(\mathbf{z})$ turns $\mathbf{z} \in \mathbb{R}^d$ into a diagonal matrix of size $d \times d$ that has the values of \mathbf{z} on its main diagonal. We denote the transpose of a matrix \mathbf{A} as \mathbf{A}' .

2.1. Static Graphs

A (*static*) graph is represented as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where $\mathcal{V} = \{v_1, v_2, \dots, v_{|\mathcal{V}|}\}$ is the set of vertices and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of edges. Vertices are also called *nodes* and we use the two terms interchangeably. Edges are also called *links* and we use the two terms interchangeably.

Several matrices can be associated with a graph. An *adjacency matrix* $\mathbf{A} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ is a matrix where $\mathbf{A}[i][j] = 0$ if $(v_i, v_j) \notin \mathcal{E}$; otherwise $\mathbf{A}[i][j] \in \mathbb{R}_+$ represents the weight of the edge. For unweighted graphs, all non-zero $\mathbf{A}[i][j]$ s are 1. A *degree matrix* $\mathbf{D} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ is a diagonal matrix where $\mathbf{D}[i][i] = \sum_{j=1}^{|\mathcal{V}|} \mathbf{A}[i][j]$ represents the degree of v_i . A *graph Laplacian* is defined as $\mathbf{L} = \mathbf{D} - \mathbf{A}$.

A graph is *undirected* if the order of the nodes in the edges is *not* important. For an undirected graph, the adjacency matrix is symmetric, i.e. $\mathbf{A}[i][j] = \mathbf{A}[j][i]$ for all i and j (in other words, $\mathbf{A} = \mathbf{A}'$). A graph is *directed* if the order of the nodes in the edges is important. Directed graphs are also called *digraphs*. For an edge (v_i, v_j) in a digraph, we call v_i the *source* and v_j the *target* of the edge. A graph is *bipartite* if the nodes can be split into two groups where there is no edge between any pair of nodes in the same group. A *multigraph* is a graph where multiple edges can exist between two nodes. A graph is *attributed* if each node is associated with some properties representing its characteristics. For a node v in an attributed graph, we let \mathbf{x}_v represent the attribute values of v . When all nodes have the same attributes, we represent all attribute values of the nodes by a matrix \mathbf{X} whose i^{th} row corresponds to the attribute values of v_i .

A *knowledge graph (KG)* corresponds to a multi-digraph with labeled edges, where the labels represent the types of the relationships. Let $\mathcal{R} = \{r_1, r_2, \dots, r_{|\mathcal{R}|}\}$ be a set of relation types. Then $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{R} \times \mathcal{V}$. That is, each edge is a triple of the form (*source, relation, target*). A KG can be attributed in which case each node $v \in \mathcal{V}$ is associated with a vector \mathbf{x}_v of attribute values. A digraph is a special case of a KG with only one relation type. An undirected graph is a special case of a KG with only one symmetric relation type.

Closely related to KGs are heterogeneous information networks. A *heterogeneous information network (HIN)* is typically defined as a digraph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with two additional functions: one mapping each node $v \in \mathcal{V}$ to a node type and one mapping each edge $(v_i, v_j) \in \mathcal{E}$ to an edge type (Shi et al. (2016); Sun and Han (2013)). Compared to KGs, HINs define node types explicitly using a mapping function whereas KGs typically define node types using triples, e.g., (*Zootopia, type, Movie*). Moreover, the definition of HINs implies the possibility of only one edge between two nodes whereas KGs allow multiple edges with different labels. However, other definitions have been considered for HINs which allow multiple edges between




Figure 1: Four graphs to be used as running examples throughout the survey. (a) and (b) are two examples of undirected graphs. They can be also considered as two snapshots of a discrete-time dynamic graph. (c) is an example of a knowledge graph. (d) is an example of a continuous-time dynamic graph where the only possible event/observation is edge addition.

two entities as well (see, e.g., Yang et al. (2012)). Despite slight differences in definition, the terms KG and HIN have been used interchangeably in some works (see, e.g., Nickel et al. (2016a)). In this work, we mainly adopt the term KG.

Example 1 Figure 1(a) represents an undirected graph with three nodes v_1 , v_2 and v_3 and three edges (v_1, v_2) , (v_1, v_3) and (v_2, v_3) . Figure 1(b) represents a graph with four nodes and four edges. The adjacency, degree, and Laplacian matrices for the graph in Figure 1(b) are as follows:

$$\mathbf{A} = \begin{bmatrix} 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \quad \mathbf{D} = \begin{bmatrix} 2 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad \mathbf{L} = \begin{bmatrix} 2 & -1 & -1 & 0 \\ -1 & 3 & -1 & -1 \\ -1 & -1 & 2 & 0 \\ 0 & -1 & 0 & 1 \end{bmatrix}$$

where the i^{th} row (and the i^{th} column) corresponds to v_i . Since the graph is undirected, \mathbf{A} is symmetric. Figure 1(c) represents a KG with four nodes v_1 , v_2 , v_3 and v_4 , three relation types r_1 , r_2 , and r_3 , and five labeled edges as follows:

$$(v_1, r_1, v_2) \quad (v_1, r_1, v_3) \quad (v_1, r_2, v_3) \quad (v_2, r_3, v_4) \quad (v_4, r_3, v_2)$$

The KG in Figure 1(c) is directed and is a multigraph as there are, e.g., two edges (with the same direction) between v_1 and v_3 .

2.2. Dynamic Graphs

We represent a *continuous-time dynamic graph (CTDG)* as a pair $(\mathcal{G}, \mathcal{O})$ where \mathcal{G} is a static graph representing an initial state of a dynamic graph at time t_0 and \mathcal{O} is a set of observations/events where each observation is a tuple of the form $(event\ type, event, timestamp)$. An event type can be an edge addition, edge deletion, node addition, node deletion, node splitting, node merging, etc. At any point $t \geq t_0$ in time, a snapshot \mathcal{G}^t (corresponding to a static graph) can be obtained from a CTDG by updating \mathcal{G} sequentially according to the observations \mathcal{O} that occurred before (or at) time t (sometimes, the update may require aggregation to handle multiple edges between two nodes).

A *discrete-time dynamic graph (DTDG)* is a sequence of snapshots from a dynamic graph sampled at regularly-spaced times. Formally, we define a DTDG as a set $\{\mathcal{G}^1, \mathcal{G}^2, \dots, \mathcal{G}^T\}$ where $\mathcal{G}^t = \{\mathcal{V}^t, \mathcal{E}^t\}$ is the graph at snapshot t , \mathcal{V}^t is the set of nodes in \mathcal{G}^t , and \mathcal{E}^t is the set of edges in \mathcal{G}^t . We use the term *dynamic graph* to refer to both DTDGs and CTDGs. Compared to a CTDG, a DTDG may lose information by looking only at some snapshots of the graph over time, but developing models for DTDGs may be generally easier. In particular, a model developed for CTDGs may be used for DTDGs, but the reverse is not necessarily true.

An *undirected dynamic graph* is a dynamic graph where at any time t , \mathcal{G}^t is an undirected graph. A *directed dynamic graph* is a dynamic graph where at any time t , \mathcal{G}^t is a digraph. A *bipartite dynamic graph* is a dynamic graph where at any time t , \mathcal{G}^t is a bipartite graph. A *dynamic KG* is a dynamic graph where at any time t , \mathcal{G}^t is a KG.

Example 2 Consider a CTDG as $(\mathcal{G}, \mathcal{O})$ where \mathcal{G} is a graph with five nodes v_1, v_2, v_3, v_4 and v_5 and with no edges between any pairs of nodes, and \mathcal{O} is:

$$\{(AddEdge, (v_2, v_5), t_1), (AddEdge, (v_1, v_2), t_2), (AddEdge, (v_1, v_4), t_3), \\ (AddEdge, (v_2, v_4), t_4), (AddEdge, (v_3, v_4), t_5)\}$$

This CTDG may be represented graphically as in Figure 1(d). The only type of observation in this CTDG is the addition of new edges. The second element of each observation corresponding to an edge addition represents the source and the target nodes of the new edge. The third element of each observation represents the timestamp at which the observation was made.

Example 3 Consider an undirected CTDG whose initial state is as in Figure 1(a). Suppose \mathcal{O} is:

$$\{(AddNode, v_4, t_1), (AddEdge, (v_2, v_4), t_2)\}$$

where $t_1 \leq t_2$. Now consider a DTDG that takes two snapshots from this CTDG, one snapshot at time t_0 and one snapshot at time t_2 . The two snapshots of this DTDG look like the graphs in Figure 1(a) and Figure 1(b) respectively.

2.3. Prediction Problems

In this survey, we mainly study three general problems for dynamic graphs: *node classification*, *edge prediction*, and *graph classification*. Node classification is the problem of classifying each node into one class from a set of predefined classes. Link prediction is the problem of predicting new links between the nodes. Graph classification is the problem of classifying a whole graph into one class from a set of predefined classes. A high-level description of some other prediction problems can be found in Section 7.1.

Reasoning over dynamic graphs typically falls under two settings: *interpolation* and *extrapolation*. Consider a dynamic graph that has incomplete information from the time interval $[t_0, t_T]$. The *interpolation* problem is to make predictions at some time t such that $t_0 \leq t \leq t_T$. The interpolation problem is also known as the *completion* problem and is mainly used for completing (dynamic) KGs (Jiang et al. (2016); Leblay and Chekol (2018);

Symbols and abbreviations	Meaning
DTDG, CTDG	Discrete-Time and Continuous-Time Dynamic Graph
KG	Knowledge Graph
$\mathcal{G}, \mathcal{V}, \mathcal{E}$	Graph, nodes, and edges.
$\mathbf{A}, \mathbf{L}, \mathbf{D}, \mathbf{X}$	Adjacency, Laplacian, degree, and attribute matrices of a graph
\mathcal{O}	Set of observations for a CTDG
\mathbf{W}	Matrix of learnable weights
$\mathcal{G}^t, \mathcal{V}^t, \mathcal{E}^t, \mathbf{A}^t$	Graph, nodes, edges, and adjacency matrix at time t .
\mathbf{v}, \mathbf{u}	Two generic nodes in a graph.
T	The number of snapshots in a DTDG
EMB	The embedding function
$[\mathbf{z}_1; \mathbf{z}_2]$	Concatenation of two vectors \mathbf{z}_1 and \mathbf{z}_2
ϕ, σ	A generic and the Sigmoid activation function
$\text{vec}(\cdot)$	Vectorized view of the input matrix or tensor
$\ \mathbf{z}\ _i, \ \mathbf{Z}\ _F$	Norm i of \mathbf{z} , and Frobenius norm of \mathbf{Z} .
\mathbf{A}', \mathbf{z}'	Transpose of a matrix and a vector

Table 1: Summary of the main notation and abbreviations.

García-Durán et al. (2018); Dasgupta et al. (2018); Goel et al. (2020)). An example of the interpolation problem is to predict which country won the world cup in 2002 assuming this information is missing in the KG. The *extrapolation* problem is to make predictions at time t such that $t \geq t_T$, i.e., predicting future based on the past. Extrapolation is usually a more challenging problem than the interpolation problem. An example of the extrapolation problem is predicting which country will win the next world cup.

Streaming scenario: In the streaming scenario, new observations are being streamed to the model at a fast rate and the model needs to update itself based on these observations in real-time so it can make informed predictions immediately after each observation arrives. For this scenario, a model may not have enough time to retrain completely or in part when new observations arrive. Streaming scenarios are often best handled by CTDGs and often give rise to extrapolation problems.

2.4. The Encoder-Decoder Framework

Following Hamilton et al. (2017b), to deal with the large notational and methodological diversity of the existing approaches and to put the various methods on an equal notational and conceptual footing, we develop an encoder-decoder framework for dynamic graphs. Before describing the encoder-decoder framework, we define one of the main components in this framework known as *embedding*.

Definition 1 *An embedding is a function that maps every node $\mathbf{v} \in \mathcal{V}$ of a graph, and every relation type $r \in \mathcal{R}$ in case of a KG, to a hidden representation where the hidden representation is typically a tuple of one or more scalars, vectors, and/or matrices of numbers. The vectors and matrices in the tuple are supposed to contain the necessary information about the nodes and relations to enable making predictions about them.*

For each node \mathbf{v} and relation r , we refer to the hidden representation of \mathbf{v} and r as the embedding of \mathbf{v} and the embedding of r respectively. When the main goal is link prediction, some works define the embedding function as mapping each pair of nodes into a hidden

representation. In these cases, we refer to the hidden representation of a pair (v, u) of nodes as the embedding of the pair (v, u) .

Having the above definition, we can now formally define an encoder and a decoder.

Definition 2 *An encoder takes as input a dynamic graph and outputs an embedding function that maps nodes, and relations in case of a KG, to hidden representations.*

Definition 3 *A decoder takes as input an embedding function and makes predictions (such as node classification, edge prediction, etc.) based on the embedding function.*

In many cases (e.g., Kipf and Welling (2017); Hamilton et al. (2017a); Yang et al. (2015); Bordes et al. (2013); Nickel et al. (2016b); Dong et al. (2014)), the embedding function $\text{EMB}(\cdot)$ maps each node, and each relation in the case of a KG, to a tuple containing a single vector; that is $\text{EMB}(v) = (\mathbf{z}_v)$ where $\mathbf{z}_v \in \mathbb{R}^{d_1}$ and $\text{EMB}(r) = (\mathbf{z}_r)$ where $\mathbf{z}_r \in \mathbb{R}^{d_2}$. Other works consider different representations. For instance, Kazemi and Poole (2018c) define $\text{EMB}(v) = (\mathbf{z}_v, \bar{\mathbf{z}}_v)$ and $\text{EMB}(r) = (\mathbf{z}_r, \bar{\mathbf{z}}_r)$, i.e. mapping each node and each relation to two vectors where each vector has a different usage. Nguyen et al. (2016) define $\text{EMB}(v) = (\mathbf{z}_v)$ and $\text{EMB}(r) = (\mathbf{z}_r, \mathbf{P}_r, \mathbf{Q}_r)$, i.e. mapping each node to a single vector but mapping each relation to a vector and two matrices. We will describe these approaches (and many others) in the upcoming sections.

A *model* corresponds to an encoder-decoder pair. One of the benefits of describing models in an encoder-decoder framework is that it allows for creating new models by combining the encoder from one model with the decoder from another model when the hidden representations produced by the encoder conform to the hidden representations consumed by the decoder.

2.4.1. TRAINING

For many choices of an encoder-decoder pair, it is possible to train the two components end-to-end. In such cases, the parameters of the encoder and the decoder are typically initialized randomly. Then, until some criterion is met, several epochs of stochastic gradient descent are performed where in each epoch, the embedding function is produced by the encoder, predictions are made based on the embedding function by the decoder, the error in predictions is computed with respect to a loss function, and the parameters of the model are updated based on the loss.

For node classification and graph classification, the loss function can be any classification loss (e.g., cross-entropy loss). For link prediction, typically one only has access to positive examples corresponding to the links already in the graph. A common approach in such cases is to generate a set of negative samples where negative samples correspond to edges that are believed to have a low probability of being in the graph. Then, having a set of positive and a set of negative samples, the training of a link predictor turns into a classification problem and any classification loss can be used. The choice of the loss function depends on the application.

2.5. Expressivity

The expressivity of the models for (dynamic) graphs can be thought of as the diversity of the graphs they can represent. Depending on the problem at hand (e.g., node classification,

link prediction, graph classification, etc.), the expressivity can be defined differently. We first provide some intuition on the importance of expressivity using the following example.

Example 4 Consider a model \mathcal{M} for binary classification (with labels *True* and *False*) in KGs. Suppose the encoder of \mathcal{M} maps every node to a tuple containing a single scalar representing the number of incoming edges to that node (regardless of the labels of the edges). For the KG in Figure 1(c), for instance, this encoder will output an embedding function as:

$$\text{EMB}(\mathbf{v}_1) = (0) \quad \text{EMB}(\mathbf{v}_2) = (2) \quad \text{EMB}(\mathbf{v}_3) = (2) \quad \text{EMB}(\mathbf{v}_4) = (1)$$

No matter what the decoder of \mathcal{M} is, Since $\text{EMB}(\mathbf{v}_2)$ and $\text{EMB}(\mathbf{v}_3)$ are identical, any deterministic decoder will assign the same class to \mathbf{v}_2 and \mathbf{v}_3 . Therefore, no matter what decoder \mathcal{M} uses, \mathcal{M} is not expressive enough to assign different classes to \mathbf{v}_2 and \mathbf{v}_3 .

From Example 4, we can see why the expressivity of a model may be important. A model that is not expressive enough is doomed to underfitting. Expressivity of the representation learning models for graphs has been the focus of several studies. It has been studied from different perspectives and for different classes of models. Xu et al. (2019b), Morris et al. (2019), Maron et al. (2019), Keriven and Peyré (2019) and Chen et al. (2019c) study the expressivity of a class of models called graph convolutional networks (see Section 3.1.6). Kazemi and Poole (2018c), Trouillon et al. (2017), Fatemi et al. (2019a), Balažević et al. (2019) and several other works provide expressivity results for models operating on KGs (see Section 3.2.2). Goel et al. (2020) provide expressivity results for their model developed for temporal KGs (see Section 5.2). We will refer to several of these works in the next sections when describing different (classes of) models.

In what follows, we provide general definitions for the expressivity of representation learning models for graphs. Before giving the definitions, we describe symmetric nodes. Two nodes in a graph are *symmetric* if there exists the same information about them (i.e. they have the same attribute values and the same neighbors). Recall that a model corresponds to an encoder-decoder pair.

Definition 4 A model \mathcal{M} with parameters Θ is fully expressive with respect to node classification if given any graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and any function $\Omega : \mathcal{V} \rightarrow \mathcal{C}$ mapping nodes to classes (where symmetric nodes are mapped to the same class), there exists an instantiation of Θ such that \mathcal{M} classifies the nodes in \mathcal{V} according to Ω .

Example 5 Consider a model \mathcal{M} for binary classification (with labels *True* and *False*) whose encoder is the one introduced in Example 4 and whose decoder is a logistic regression model. It can be verified that the encoder has no parameters so model parameters Θ correspond to the parameters of the decoder (i.e. the logistic regression). We disprove the full expressivity of \mathcal{M} using a counterexample. According to Definition 4, a counterexample corresponds to a pair $(\mathcal{G}_{ce}, \Omega_{ce})$ (where *ce* stands for counterexample) of a specific graph and a specific function Ω_{ce} such that there exists no instantiation of Θ that classifies the nodes of \mathcal{G}_{ce} according to Ω_{ce} . Let \mathcal{G}_{ce} be the graph in Figure 1(c) and let Ω_{ce} be a mapping function defined as: $\Omega_{ce}(\mathbf{v}_1) = \text{True}$, $\Omega_{ce}(\mathbf{v}_2) = \text{True}$, $\Omega_{ce}(\mathbf{v}_3) = \text{False}$, and $\Omega_{ce}(\mathbf{v}_4) = \text{True}$. From Example 4, we know that the encoder gives the same embedding for \mathbf{v}_2 and \mathbf{v}_3 so there cannot exist an instantiation of Θ which classifies \mathbf{v}_2 as *True* and \mathbf{v}_3 as *False*. Hence, the above pair is a counterexample.

A similar definition can be given for the full expressivity of a model with respect to link prediction and graph classification.

Definition 5 *A model \mathcal{M} with parameters Θ is fully expressive with respect to link prediction if given any graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and any function $\Omega : \mathcal{E} \rightarrow \{\text{True}, \text{False}\}$ indicating the existence or non-existence of (labeled) edges for all node-pairs in the graph, there exists an instantiation of Θ such that \mathcal{M} classifies the edges in \mathcal{E} according to Ω .*

Definition 6 *A model \mathcal{M} with parameters Θ is fully expressive with respect to graph classification if given any set $\mathcal{S} = \{\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_n\}$ of non-isomorphic graphs and any function $\Omega : \mathcal{S} \rightarrow \mathcal{C}$ mapping graphs to classes, there exists an instantiation of Θ such that \mathcal{M} classifies the graphs according to Ω .*

2.6. Sequence Models

In dynamic environments, data often consists of sequences of observations of varying lengths. There is a long history of models to handle sequential data without a fixed length. This includes auto-regressive models (Akaike (1969)) that predict the next observations based on a window of past observations. Alternatively, since it is not always clear how long the window of past observations should be, hidden Markov models (Rabiner and Juang (1986)), Kalman filters (Welch et al. (1995)), dynamic Bayesian networks (Murphy and Russell (2002)) and dynamic conditional random fields (Sutton et al. (2007)) use hidden states to capture relevant information that might be arbitrarily far in the past. Today, those models can be seen as special cases of recurrent neural networks, which allow rich and complex hidden dynamics.

Recurrent neural networks (RNNs) (Elman (1990); Cho et al. (2014)) have achieved impressive results on a range of sequence modeling problems. The core principle of the RNN is that its input is a function of the current data point as well as the history of the previous inputs. A simple RNN model can be formulated as follows:

$$\mathbf{h}^t = \phi(\mathbf{W}_i \mathbf{x}^t + \mathbf{W}_h \mathbf{h}^{t-1} + \mathbf{b}_i) \quad (1)$$

where $\mathbf{x}^t \in \mathbb{R}^{d_{in}}$ is the input at position t in the sequence, $\mathbf{h}^{t-1} \in \mathbb{R}^d$ is a hidden representation containing information about the sequence of inputs until time $t - 1$, $\mathbf{W}_i \in \mathbb{R}^{d \times d_{in}}$ and $\mathbf{W}_h \in \mathbb{R}^{d \times d}$ are weight matrices, $\mathbf{b}_i \in \mathbb{R}^d$ represents the vector of biases, ϕ is an activation function, and $\mathbf{h}^t \in \mathbb{R}^d$ is an updated hidden representation containing information about the sequence of inputs until time t . We use $\mathbf{h}^t = \text{RNN}(\mathbf{h}^{t-1}, \mathbf{x}^t)$ to represent the output of an RNN operation on a previous state \mathbf{h}^{t-1} and a new input \mathbf{x}^t .

Long short term memory (LSTM) (Hochreiter and Schmidhuber (1997)) is considered one of the most successful RNN architectures. The original LSTM model can be neatly defined with the following equations:

$$\mathbf{i}^t = \sigma(\mathbf{W}_{ii} \mathbf{x}^t + \mathbf{W}_{ih} \mathbf{h}^{t-1} + \mathbf{b}_i) \quad (2)$$

$$\mathbf{f}^t = \sigma(\mathbf{W}_{fi} \mathbf{x}^t + \mathbf{W}_{fh} \mathbf{h}^{t-1} + \mathbf{b}_f) \quad (3)$$

$$\mathbf{c}^t = \mathbf{f}^t \odot \mathbf{c}^{t-1} + \mathbf{i}^t \odot \text{Tanh}(\mathbf{W}_{ci} \mathbf{x}^t + \mathbf{W}_{ch} \mathbf{h}^{t-1} + \mathbf{b}_c) \quad (4)$$

$$\mathbf{o}^t = \sigma(\mathbf{W}_{oi} \mathbf{x}^t + \mathbf{W}_{oh} \mathbf{h}^{t-1} + \mathbf{b}_o) \quad (5)$$

$$\mathbf{h}^t = \mathbf{o}^t \odot \text{Tanh}(\mathbf{c}^t) \quad (6)$$

Here \mathbf{i}^t , \mathbf{f}^t , and \mathbf{o}^t represent the input, forget and output gates respectively, while \mathbf{c}^t is the memory cell and \mathbf{h}^t is the hidden state. σ and Tanh represent the Sigmoid and hyperbolic tangent activation functions respectively. Gated recurrent units (GRUs) (Cho et al. (2014)) is another successful RNN architecture. In the context of dynamic graphs, sequence models such as LSTMs and GRUs can be used to, e.g., provide node representations based on the history of the node (see Sections 4.6.1 and 4.6.3).

Fully attentive models have recently demonstrated on-par or superior performance compared to RNN variants for a variety of tasks (see, e.g., Vaswani et al. (2017); Dehghani et al. (2018); Krantz and Kalita (2018); Shaw et al. (2018)). These models rely only on (self-)attention and abstain from using recurrence. Let $\mathbf{X}_{\text{in}} \in \mathbb{R}^{T \times d}$ represent a sequence containing T elements each with d features. The idea behind a self-attention layer is to update each row of \mathbf{X}_{in} by allowing it to attend to itself and all other rows. For this purpose, Vaswani et al. (2017) first create $\bar{\mathbf{X}}_{\text{in}} = \mathbf{X}_{\text{in}} + \mathbf{P}$ where $\mathbf{P} \in \mathbb{R}^{T \times d}$ is called the *positional encoding* matrix and carries information about the position of each element in the sequence. Then, they project the matrix $\bar{\mathbf{X}}_{\text{in}}$ into a matrix $\mathbf{Q} = \bar{\mathbf{X}}_{\text{in}} \mathbf{W}_Q \in \mathbb{R}^{T \times d_k}$ dubbed *queries matrix*, a matrix $\mathbf{K} = \bar{\mathbf{X}}_{\text{in}} \mathbf{W}_K \in \mathbb{R}^{T \times d_k}$ dubbed *keys matrix*, and a matrix $\mathbf{V} = \bar{\mathbf{X}}_{\text{in}} \mathbf{W}_V \in \mathbb{R}^{T \times d_v}$ dubbed *values matrix*, where $\mathbf{W}_Q, \mathbf{W}_K \in \mathbb{R}^{d \times d_k}$ and $\mathbf{W}_V \in \mathbb{R}^{d \times d_v}$ are matrices with learnable parameters. Then, each row of the matrix $\bar{\mathbf{X}}_{\text{in}}$ is updated by taking a weighted sum of the rows in \mathbf{V} . The weights are computed using the query and key matrices. The updated matrix $\mathbf{X}_{\text{out}} \in \mathbb{R}^{T \times d_v}$ is computed as follows:

$$\mathbf{X}_{\text{out}} = \text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}'}{\sqrt{d_k}}\right)\mathbf{V} \quad (7)$$

where softmax performs a row-wise normalization of the input matrix and $\text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}'}{\sqrt{d_k}}\right)$ gives the weights. A mask can be added to Equation (7) to make sure that at time t , the mechanism only allows a sequence model to attend to the points before time t . Vaswani et al. (2017) also define a *multi-head* self-attention mechanism by considering multiple self-attention blocks (as defined in Equation (7)) each having different weight matrices and then concatenating the results. In the context of static graphs, the initial \mathbf{X}_{in} may correspond to the representations of the neighbors of a node, and in the context of dynamic graphs where node representations keep evolving, the initial \mathbf{X}_{in} may correspond to a node's representations at different points in time (see Sections 3.1.6 and 4.6.2).

2.7. Temporal Point Processes

Temporal point processes (TPPs) (Cox and Lewis (1972)) are stochastic processes which are used for modeling sequential asynchronous discrete events occurring in continuous time. A typical realization of a TPP is a sequence of discrete events occurring at time points t_1, t_2, t_3, \dots for $t_i \leq T$, where T represents the time horizon of the process. A TPP is generally characterized using a conditional intensity function $\lambda(t)$ such that $\lambda(t)dt$ represents the probability of an event happening in the interval $[t, t + dt]$ given the history t_1, \dots, t_n of the process and given that no event occurred until $t_n < t \leq T$. The conditional density function $\mathbf{f}(t)$, indicating the density of the occurrence of the next event at some time point $t_n < t \leq T$, can be obtained as $\mathbf{f}(t) = \lambda(t)\mathbf{S}(t)$. Here, $\mathbf{S}(t) = \exp\left(-\int_{t_n}^t \lambda(\tau)d\tau\right)$, called the

survival function of the process, is the probability that no event happens during $[t_n, t)$. The time for the next event can be predicted by taking an expectation over $\mathbf{f}(t)$.

Traditionally, intensity functions were hand-designed to model how future/present events depend on the past events in the TPP. Some of the well-known TPPs include Hawkes process (Hawkes (1971); Mei and Eisner (2017)), Poisson processes (Kingman (2005)), self-correcting processes (Isham and Westcott (1979)), and autoregressive conditional duration processes (Engle and Russell (1998)). Depending on the application, one may use the intensity function in one of these TPPs or design new ones. Recently, there has been growing interest in learning the intensity function entirely from the data (see, e.g., Du et al. (2016)). In the context of dynamic graphs, a TPP with an intensity function parameterized by the node representations in the graph can be constructed to predict when something will happen to a single node or to a pair of nodes (see Section 5.1).

3. Representation Learning for Static Graphs

In this section, we provide an overview of representation learning approaches for static graphs. The main purpose of this section is to provide enough information for the descriptions and discussions in the next sections on dynamic graphs. Readers interested in learning more about representation learning on static graphs can refer to several existing surveys specifically written on this topic (e.g., see Hamilton et al. (2017b); Zhang et al. (2018b); Cai et al. (2018); Cui et al. (2018) for graphs and Nickel et al. (2016a); Wang et al. (2017a) for KGs).

3.1. Encoders

As described in Subsection 2.4, a model can be viewed as a combination of an encoder and a decoder. In this section, we describe different approaches for creating encoders.

3.1.1. HIGH-ORDER PROXIMITY MATRICES

While the adjacency matrix of a graph only represents local proximities, one can also define *high-order proximity* matrices (Ou et al. (2016)) also known as *graph similarity metrics* (da Silva Soares and Prudêncio (2012)). Let \mathbf{S} be a high-order proximity matrix. A simple approach for creating an encoder is to let $\text{EMB}(\mathbf{v}_i) = (\mathbf{S}[i])$ (or $\text{EMB}(\mathbf{v}_i) = (\mathbf{S}'[i])$) corresponding to the i^{th} row (or the i^{th} column) of matrix \mathbf{S} . Encoders based on high-order proximity matrices are typically parameter-free and do not require learning (although some of them have hyper-parameters that need to be tuned). In what follows, we describe several of these matrices.

Common neighbors matrix is defined as $\mathbf{S}_{CN} = \mathbf{A}\mathbf{A}$. $\mathbf{S}_{CN}[i][j]$ corresponds to the number of nodes that are connected to both \mathbf{v}_i and \mathbf{v}_j . For a directed graph, $\mathbf{S}_{CN}[i][j]$ counts how many nodes \mathbf{v} are simultaneously the target of an edge starting at \mathbf{v}_i and the source of an edge ending at \mathbf{v}_j .

Jaccard's coefficient is a slight modification of \mathbf{S}_{CN} where one divides the number of common neighbors of \mathbf{v}_i and \mathbf{v}_j by the total number of distinct nodes that are the targets of edges starting at \mathbf{v}_i or the sources of edges ending at \mathbf{v}_j . Formally, Jaccard's coefficient is defined as $\mathbf{S}_{JC}[i][j] = \mathbf{S}_{CN}[i][j] / (\sum_{k=1}^{|\mathcal{V}|} (\mathbf{A}[i][k] + \mathbf{A}[k][j]) - \mathbf{S}_{CN}[i][j])$.

Adamic-Adar is defined as $\mathbf{S}_{AA} = \mathbf{A}\hat{\mathbf{D}}\mathbf{A}$, where $\hat{\mathbf{D}}[i][i] = 1/\sum_{k=1}^{|\mathcal{V}|}(\mathbf{A}[i][k] + \mathbf{A}[k][i])$. \mathbf{S}_{AA} computes the weighted sum of common neighbors where the weight is inversely proportional to the degree of the neighbor.

Katz index is defined as $\mathbf{S}_{Katz} = \sum_{k=1}^{\infty}(\beta\mathbf{A})^k$. $\mathbf{S}_{Katz}[i][j]$ corresponds to a weighted sum of all the paths between two nodes v_i and v_j . β controls the depth of the connections: the closer β is to 1, the longer paths one wants to consider. One can rewrite the formula recursively as $\beta\mathbf{A}\mathbf{S}_{Katz} + \beta\mathbf{A} = \mathbf{S}_{Katz}$ and, as a corollary, obtain $\mathbf{S}_{Katz} = (\mathbf{I}_N - \beta\mathbf{A})^{-1}\beta\mathbf{A}$.

Preferential Attachment is simply a product of in- and out- degrees of nodes: $\mathbf{S}_{PA}[i][j] = (\sum_{k=1}^{|\mathcal{V}|} \mathbf{A}[i][k])(\sum_{k=1}^{|\mathcal{V}|} \mathbf{A}[k][j])$.

3.1.2. SHALLOW ENCODERS

Shallow encoders first decide on the number and shape of the vectors and matrices for node and relation embeddings. Then, they consider each element in these vectors and matrices as a parameter to be directly learned from the data. A shallow encoder can be viewed as a lookup function that finds the hidden representation corresponding to a node or a relation given their id. Shallow encoders are commonly used for KG embedding (see e.g., Nickel et al. (2011); Yang et al. (2015); Trouillon et al. (2016); Bordes et al. (2013); Nguyen et al. (2016); Kazemi and Poole (2018c); Dettmers et al. (2018)).

3.1.3. DECOMPOSITION APPROACHES

Decomposition methods are among the earliest attempts for developing encoders for graphs. They learn node embeddings similar to shallow encoders but in an unsupervised way: the node embeddings are learned in a way that connected nodes are close to each other in the embedded space. Once the embeddings are learned, they can be used for purposes other than reconstructing the edges (e.g., for clustering). Formally, for an undirected graph \mathcal{G} , learning node embeddings $\text{EMB}(v_i) = (\mathbf{z}_{v_i})$, where $\mathbf{z}_{v_i} \in \mathbb{R}^d$, such that connected nodes are close in the embedded space can be done through solving the following optimization problem:

$$\min_{\{\mathbf{z}_{v_i}\}_{i=1}^N} \sum_{i,j} \mathbf{A}[i][j] \|\mathbf{z}_{v_i} - \mathbf{z}_{v_j}\|^2 \tag{8}$$

This loss ensures that connected nodes are close to each other in the embedded space. One needs to impose some constraints to get rid of a scaling factor and to eliminate the trivial solution where all nodes are set to a single vector. For that let us consider a new matrix $\mathbf{Y} \in \mathbb{R}^{|\mathcal{V}| \times d}$, such that its rows give the embedding: $\mathbf{Y}[i] = \mathbf{z}'_{v_i}$. Then one can add the constraints to the optimization problem (8): $\mathbf{Y}'\mathbf{D}\mathbf{Y} = \mathbf{I}$, where \mathbf{D} is a diagonal matrix of degrees as defined in Subsection 2.1. As was proved by Belkin and Niyogi (2001), this constrained optimization is equivalent to solving a generalized eigenvalue decomposition:

$$\mathbf{L}\mathbf{y} = \lambda\mathbf{D}\mathbf{y}, \tag{9}$$

where \mathbf{L} is a graph Laplacian; the matrix \mathbf{Y} can be obtained by considering the $|\mathcal{V}| \times d$ matrix of top- d generalized eigenvectors: $\mathbf{Y} = [\mathbf{y}_1 \dots \mathbf{y}_d]$.

Sussman et al. (2012) suggested to use a slightly different embedding based on the eigenvalue decomposition of the adjacency matrix $\mathbf{A} = \mathbf{U}\Sigma\mathbf{U}'$ (this matrix is symmetric

for an undirected graph). Then one can choose the top d eigenvalues $\{\lambda_1, \dots, \lambda_d\}$ and the corresponding eigenvectors $\{\mathbf{u}_1, \dots, \mathbf{u}_d\}$ and construct a new matrix

$$\mathbf{Z} = \mathbf{U}_{\leq d} \sqrt{\boldsymbol{\Sigma}_{\leq d}} \in \mathbb{R}^{|\mathcal{V}| \times d}, \quad (10)$$

where $\boldsymbol{\Sigma}_{\leq d} = \text{diag}(\lambda_1, \dots, \lambda_d)$, and $\mathbf{U}_{\leq d} = [\mathbf{u}_1 \dots \mathbf{u}_d]$. Rows of this matrix can be used as node embedding: $\mathbf{z}_{v_i} = \mathbf{Z}[i]' \in \mathbb{R}^d$. This is the so called *adjacency spectral embedding*, see also Levin et al. (2018).

For directed graphs, because of their asymmetric nature, keeping track of the n^{th} -order neighbors where $n > 1$ becomes difficult. For this reason, working with a high-order proximity matrix \mathbf{S} is preferable (see Section 3.1.1 for a description of high-order proximity matrices). Moreover, for directed graphs, it may be preferable to learn two vector representations per node, one to be used when the node is the source and the other to be used when the node is the target of an edge. One may learn embeddings for directed graphs by solving the following:

$$\min_{\mathbf{Z}_s, \mathbf{Z}_t} \|\mathbf{S} - \mathbf{Z}_s \mathbf{Z}_t'\|_F^2, \quad (11)$$

where $\|\cdot\|_F$ is the Frobenius norm and $\mathbf{Z}_s, \mathbf{Z}_t \in \mathbb{R}^{|\mathcal{V}| \times d}$. Given the solution, one can define the *source features* of a node v_i as $\mathbf{Z}_s[i]'$ and the *target features* as $\mathbf{Z}_t[i]'$. A single-vector embedding of a node v_i can be defined as a concatenation of these features. The Eckart–Young–Mirsky theorem (Eckart and Young (1936)) from linear algebra indicates that the solution is equivalent to finding the singular value decomposition of \mathbf{S} :

$$\mathbf{S} = \mathbf{U}_s \boldsymbol{\Sigma} (\mathbf{U}_t)'\!, \quad (12)$$

where $\boldsymbol{\Sigma} = \text{diag}(\sigma_1, \dots, \sigma_{|\mathcal{V}|})$ is a matrix of singular values and \mathbf{U}_s and \mathbf{U}_t are matrices of left and right singular vectors respectively (stacked as columns). Then using the top d singular vectors one gets the solution of the optimization problem in (11):

$$\mathbf{Z}_s = (\mathbf{U}_s)_{\leq d} \sqrt{\boldsymbol{\Sigma}_{\leq d}} \quad (13)$$

$$\mathbf{Z}_t = (\mathbf{U}_t)_{\leq d} \sqrt{\boldsymbol{\Sigma}_{\leq d}}. \quad (14)$$

3.1.4. RANDOM WALK APPROACHES

A popular class of approaches for learning an embedding function for graphs is the class of random walk approaches. Similar to decomposition approaches, encoders based on random walks also learn embeddings in an unsupervised way. However, compared to decomposition approaches, these embeddings may capture longer-term dependencies. To describe the encoders in this category, first we define what a random walk is and then describe the encoders that leverage random walks to learn an embedding function.

Definition 7 *A random walk for a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a sequence of nodes v_1, v_2, \dots, v_l where $v_i \in \mathcal{V}$ for all $1 \leq i \leq l$ and $(v_i, v_{i+1}) \in \mathcal{E}$ for all $1 \leq i \leq l - 1$. l is called the length of the walk.*

A random walk of length l can be generated by starting at a node v_i in the graph, then transitioning to a neighbor v_j of v_i ($j \neq i$), then transitioning to a neighbor of v_j and continuing this process for l steps. The selection of the first node and the node to transition to in each step can be uniformly at random or based on some distribution/strategy.

Example 6 Consider the graph in Figure 1(b). The following are three examples of random walks on this graph with length 4: 1) v_1, v_3, v_2, v_3 , 2) v_2, v_1, v_2, v_4 and 3) v_4, v_2, v_4, v_2 . In the first walk, the initial node has been selected to be v_1 . Then a transition has been made to v_3 , which is a neighbor of v_1 . Then a transition has been made to v_2 , which is a neighbor of v_3 and then a transition back to v_3 , which is a neighbor of v_2 . The following are two examples of invalid random walks: 1) v_1, v_4, v_2, v_3 and 2) v_1, v_3, v_4, v_2 . The first one is not a valid random walk since a transition has been made from v_1 to v_4 when there is no edge between v_1 and v_4 . The second one is not valid because a transition has been made from v_3 to v_4 when there is no edge between v_3 and v_4 .

Random walk encoders perform multiple random walks of length l on a graph and consider each walk as a sentence, where the nodes are considered as the words of these sentences. Then they use the techniques from natural language processing for learning word embeddings (e.g., Mikolov et al. (2013); Pennington et al. (2014)) to learn a vector representation for each node in the graph. One such approach is to create a matrix \mathbf{S} from these random walks such that $\mathbf{S}[i][j]$ corresponds to the number of times v_i and v_j co-occurred in random walks and then factorize the matrix (see Section 3.1.3) to get vector representations for nodes.

Random walk encoders typically differ in the way they perform the walk, the distribution they use for selecting the initial node, and the transition distribution they use. For instance, DeepWalk (Perozzi et al. (2014)) selects both the initial node and the node to transition to uniformly at random. Perozzi et al. (2016) extends DeepWalk by allowing random walks to skip over multiple nodes at each transition. Node2Vec (Grover and Leskovec (2016)) selects the node to transition to based on a combination of breadth-first search (to capture local information) and depth-first search (to capture global information).

Random walk encoders have been extended to KGs (and HINs) through constraining the walks to conform to some meta-paths. A meta-path can be considered as a sequence of relations in \mathcal{R} . Dong et al. (2017) propose *metapath2vec* where each random walk is constrained to conform to a meta-path r_1, r_2, \dots, r_k by starting randomly at a node v_1 whose type is compatible with the source type of r_1 . Then the walk transitions to a node v_2 where v_2 is selected uniformly at random among the nodes having relation r_1 with v_1 , then the walk transitions to a node v_3 where v_3 is selected uniformly at random among the nodes having relation r_2 with v_2 , and so forth. Each meta-path provides a semantic relationship between the start and end nodes.

Shi et al. (2018) take a similar approach as *metapath2vec* but aim at learning node embeddings that are geared more towards improving recommendation performance. Both Dong et al. (2017) and Shi et al. (2018) use a set of hand-crafted meta-paths to guide the random walks. Instead of hand-crafting meta-paths, Chen and Sun (2017) propose a greedy approach to select the meta-paths based on performance on a validation set. Zhang et al. (2018a) identify some limitations for models restricting random walks to conform to meta-paths. They propose meta-graphs as an alternative to meta-paths in which relations

are connected as a graph (instead of a sequence) and at each node, the walk can select to conform to any outgoing edge in the meta-graph. Ristoski and Paulheim (2016) extend random walk approaches to general RDF data.

3.1.5. AUTOENCODER APPROACHES

Another class of models for learning an embedding function for static graphs is by using autoencoders. Similar to the decomposition approaches, these approaches are also unsupervised. However, instead of learning shallow embeddings that reconstruct the edges of a graph, the models in this category create a deep encoder that compresses a node’s neighborhood to a vector representation, which can be then used to reconstruct the node’s neighborhood. The model used for compression and reconstruction is referred to as an autoencoder. Similar to the decomposition approaches, once the node embeddings are learned, they may be used for purposes other than predicting a node’s neighborhood.

In its simplest form, an autoencoder (Hinton and Salakhutdinov (2006)) contains two components called the *encoder* and *reconstructor*¹, where each component is a feed-forward neural network. The encoder takes as input a vector $\mathbf{a} \in \mathbb{R}^N$ (e.g., corresponding to N numerical features of an object) and passes it through several feed-forward layers producing $\mathbf{z} \in \mathbb{R}^d$ such that $d \ll N$. The reconstructor receives \mathbf{z} as input and passes it through several feed-forward layers aiming at reconstructing \mathbf{a} . That is, assuming the output of the reconstructor is $\hat{\mathbf{a}}$, the two components are trained such that $\|\mathbf{a} - \hat{\mathbf{a}}\|$ is minimized. \mathbf{z} can be considered a compression of \mathbf{a} .

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a graph with adjacency matrix \mathbf{A} . For a node $v_i \in \mathcal{V}$, let $\mathbf{A}[i]$ represent the i^{th} row of the adjacency matrix corresponding to the neighbors of v_i . To use autoencoders for generating node embeddings, Wang et al. (2016) train an autoencoder (named *SDNE*) that takes a vector $\mathbf{A}[i] \in \mathbb{R}^{|\mathcal{V}|}$ as input, feeds the input vector through an encoder and produces $\mathbf{z}_i \in \mathbb{R}^d$, and then feeds \mathbf{z}_i into a reconstructor to reconstruct $\mathbf{A}[i]$. After training, the \mathbf{z}_i vectors corresponding to the output of the encoder of the autoencoder can be considered as embeddings for the nodes v_i . A graph decoder can be applied to these embeddings to make predictions. \mathbf{z}_i and \mathbf{z}_j may further be constrained to be close in Euclidean space if v_i and v_j are connected. For the case of attributed graphs, Tran (2018) concatenates the attribute values \mathbf{x}_i of node v_i to $\mathbf{A}[i]$ and feeds the concatenation $[\mathbf{x}_i; \mathbf{A}[i]]$ into an autoencoder. Cao et al. (2016) propose an autoencoder approach (named *RDNG*) that is similar to SDNE, but they first compute a high-order proximity matrix $\mathbf{S} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ based on node co-occurrences on random walks (any other matrix from Section 3.1.1 may also be used), and then feed $\mathbf{S}[i]$ s into an autoencoder.

3.1.6. GRAPH CONVOLUTIONAL NETWORK APPROACHES

Yet another class of models for learning node embeddings in a graph are graph convolutional networks (GCNs). As the name suggests, graph convolutions generalize convolutions to arbitrary graphs. Graph convolutions have spatial (see, e.g., Hamilton et al. (2017a,b); Schlichtkrull et al. (2018); Gilmer et al. (2017)) and spectral constructions (see, e.g., Liao et al. (2019); Kipf and Welling (2017); Defferrard et al. (2016); Levie et al. (2017)). Here,

1. Reconstructor is also called *decoder* but we use the name *reconstructor* to avoid confusion with graph decoders.

we describe the spatial (or message passing) view and refer the reader to Bronstein et al. (2017) for the spectral view.

A GCN consists of multiple layers where each layer takes node representations (a vector per node) as input and outputs transformed representations. Let $\mathbf{z}_{v,l}$ be the representation for a node v after passing it through the l^{th} layer. A very generic forward pass through a GCN layer transforms the representation of each node v as follows:

$$\mathbf{z}_{v,l+1} = \mathbf{transform}(\{\mathbf{z}_{v,j}\}_{0 \leq j \leq l}, \{\mathbf{z}_{u,k}\}_{u \in \mathcal{N}(v), 0 \leq k \leq l}, \Theta) \quad (15)$$

where $\mathcal{N}(v)$ represents the neighbors of v and $\mathbf{transform}$ is a function parametrized by Θ which aggregates the information from the previous representations of the neighbors of v and combines it with the previous representations of v itself to compute $\mathbf{z}_{v,l+1}$. The $\mathbf{transform}$ function should be invariant to the order of the nodes in $\mathcal{N}(v)$ because there is no specific ordering to the nodes in an arbitrary graph. Moreover, it should be able to handle a variable number of neighbors. If the graph is attributed, for each node v , $\mathbf{z}_{v,0}$ can be initialized to \mathbf{x}_v corresponding to the attribute values of v (see, e.g., Kipf and Welling (2017)). Otherwise, they can be initialized using a one-hot encoding of the nodes (see, e.g., Schlichtkrull et al. (2018)). In a GCN with L layers, each node receives information from the nodes at most L hops away from it.

There is a large literature on the design of the $\mathbf{transform}$ function (see, e.g., Li et al. (2015); Kipf and Welling (2017); Hamilton et al. (2017a); Dai et al. (2018)). Kipf and Welling (2017) formulate it as:

$$\mathbf{Z}_{l+1} = \sigma(\tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}} \mathbf{Z}_l \mathbf{W}_{l+1}) \quad (16)$$

where $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}_N$ is adjacency matrix with self-connections for input graph, N is the number of nodes in the graph, \mathbf{I}_N is the identity matrix, \mathbf{W}_{l+1} is a parameter matrix for the $(l+1)^{\text{th}}$ layer and $\sigma(\cdot)$ is a non-linearity. $\tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}} \mathbf{Z}_l$ corresponds to taking a normalized average of the features of v and its neighbors (treating the features of v and its neighbors identically). Other formulations for the $\mathbf{transform}$ function can be found in several recent surveys (see, e.g., Zhou et al. (2018a); Cai et al. (2018)).

For a node v , not all the neighboring nodes may be equally important. Veličković et al. (2018) propose an adaptive attention mechanism that learns to weigh the neighbors depending on their importance when aggregating information from the neighbors. The mechanism is adaptive in the sense that the weight of a node is not fixed and depends on the current representation of the node for which the aggregation is performed. Following Vaswani et al. (2017), Veličković et al. (2018) also use multi-headed attention. *GaAN* (Zhang et al. (2018c)) extends this idea and introduces adaptive attention weights for different attention heads, i.e., the weights for different attention heads depend on the node for which the multi-head attention is being applied.

In graphs like social networks, there can be nodes that have a large number of neighbors. This can make the $\mathbf{transform}$ function computationally prohibitive. Hamilton et al. (2017a) propose to use a uniform sampling of the neighbors to fix the neighborhood size to a constant number. Not only the sampling helps reduce computational complexity and speed up training, but also it acts as a regularizer. Ying et al. (2018a) propose an extension of this idea according to which the neighborhood of a node v is formed by repeatedly

starting truncated random walks from \mathbf{v} and choosing the nodes most frequently hit by these truncated random walks. In this way, the neighborhood of a node consists of the nodes most relevant to it, regardless of whether they are connected with an edge or not.

Expressivity: There are currently two approaches for measuring the expressivity of GCNs. Xu et al. (2019b) study the expressiveness of certain GCN models with respect to graph classification (see Definition 6) and show that in terms of distinguishing non-isomorphic graphs, these GCNs are *at most* as powerful as the Weisfeiler-Lehman isomorphism test (Weisfeiler and Lehman (1968)) — a test which is able to distinguish a broad class of graphs (Babai and Kucera (1979)) but also known to fail in some corner cases (Cai et al. (1992)). In a concurrent work, a similar result has been reported by Morris et al. (2019). Xu et al. (2019b) provide the necessary conditions under which these GCNs become as powerful as the Weisfeiler-Lehman test. On the other hand, Maron et al. (2019) and Keriven and Peyré (2019) study how well certain GCN models can approximate any continuous function which is invariant to permutation of its input. They proved that a certain class of networks, called *G-invariant*, are universal approximators. Chen et al. (2019c) demonstrate that these two approaches to the expressivity of GCNs are closely related.

GCNs for KGs: Several works extend GCNs to KG embedding. One notable example is called relational GCN (RGCN) (Schlichtkrull et al. (2018)). The core operation that RGCN does differently is the application of a relation specific transformation (i.e., the transformation depends on the direction and the label of the edge) to the neighbors of the nodes in the aggregation function. In RGCN, the **transform** function is defined as follows:

$$\mathbf{z}_{\mathbf{v},l+1} = \sigma\left(\sum_{r \in \mathcal{R}} \sum_{\mathbf{u} \in \mathcal{N}(\mathbf{v},r)} \frac{1}{c_{\mathbf{v},r}} \mathbf{W}_{r,l} \mathbf{z}_{\mathbf{u},l} + \mathbf{W}_{0,l} \mathbf{z}_{\mathbf{v},l}\right) \quad (17)$$

where \mathcal{R} is the set of relation types, $\mathcal{N}(\mathbf{v},r)$ is the set of neighboring nodes connected to \mathbf{v} via relation r , $c_{\mathbf{v},r}$ is a normalization factor that can either be learned or fixed (e.g., to $|\mathcal{N}(\mathbf{v},r)|$), $\mathbf{W}_{r,l}$ is a transformation matrix for relation r at the l^{th} layer, and $\mathbf{W}_{0,l}$ is a self-transformation matrix at the l^{th} layer. Sourek et al. (2018) and Kazemi and Poole (2018b) propose other variants for Equation (17) where (roughly) the transformations are done using soft first-order logic rules. Wang et al. (2019) and Nathani et al. (2019) propose attention-based variants of Equation (17).

3.2. Decoders

We divide the discussion on decoders into those used for graphs and those used for KGs.

3.2.1. DECODERS FOR STATIC GRAPHS

For static graphs, the embedding function usually maps each node to a single vector; that is, $\text{EMB}(\mathbf{v}) = (\mathbf{z}_{\mathbf{v}})$ where $\mathbf{z}_{\mathbf{v}} \in \mathbb{R}^d$ for any $\mathbf{v} \in \mathcal{V}$. To classify a node \mathbf{v} , a decoder can be any classifier on $\mathbf{z}_{\mathbf{v}}$ (e.g., logistic regression or random forest).

To predict a link between two nodes \mathbf{v} and \mathbf{u} , for undirected (and bipartite) graphs, the most common decoder is based on the *dot-product* of the vectors for the two nodes, i.e., $\mathbf{z}_{\mathbf{v}}' \mathbf{z}_{\mathbf{u}}$. The dot-product gives a score that can then be fed into a sigmoid function whose output can be considered as the probability of a link existing between \mathbf{v} and \mathbf{u} . Grover and Leskovec (2016) propose several other decoders for link prediction in undirected graphs.

Their decoders are based on defining a function $\mathbf{f}(\mathbf{z}_v, \mathbf{z}_u)$ that combines the two vectors \mathbf{z}_v and \mathbf{z}_u into a single vector. The resulting vector is then considered as the edge features that can be fed into a classifier. These combining functions include average $\frac{\mathbf{z}_v + \mathbf{z}_u}{2}$, Hadamard multiplication $\mathbf{z}_v \odot \mathbf{z}_u$, absolute value of the difference $\text{abs}(\mathbf{z}_v - \mathbf{z}_u)$, and squared value of the difference $(\mathbf{z}_v - \mathbf{z}_u)^2$. Instead of computing the distance between \mathbf{z}_v and \mathbf{z}_u in the Euclidean space, the distance can be computed in other spaces such as the hyperbolic space (see, e.g., Chamberlain et al. (2017)). Different spaces offer different properties. Note that all these four combination functions are symmetric, i.e., $\mathbf{f}(\mathbf{z}_v, \mathbf{z}_u) = \mathbf{f}(\mathbf{z}_u, \mathbf{z}_v)$ where \mathbf{f} is any of the above functions. This is an important property when the graph is undirected.

For link prediction in digraphs, it is important to treat the source and target of the edge differently. Towards this goal, one approach is to *concatenate* the two vectors as $[\mathbf{z}_v; \mathbf{z}_u]$ and feed the concatenation into a classifier (see, e.g., Pareja et al. (2019)). Another approach used by Ma et al. (2018b) is to project the source and target vectors to another space as $\hat{\mathbf{z}}_v = \mathbf{W}_1 \mathbf{z}_v$ and $\hat{\mathbf{z}}_u = \mathbf{W}_2 \mathbf{z}_u$, where \mathbf{W}_1 and \mathbf{W}_2 are matrices with learnable parameters, and then take the dot-product in the new space (i.e., $\hat{\mathbf{z}}_v' \hat{\mathbf{z}}_u$). A third approach is to take the vector representation \mathbf{z}_v of a node $v \in \mathcal{V}$ and send it through a feed-forward neural network with $|\mathcal{V}|$ outputs where each output gives the score for whether v has a link with one of the nodes in the graph or not. This approach is used mainly in graph autoencoders (see, e.g., Wang et al. (2016); Cao et al. (2016); Tran (2018); Goyal et al. (2017); Chen et al. (2018a)) and is used for both directed and undirected graphs.

The decoder for a graph classification task needs to compress node representations into a single representation which can then be fed into a classifier to perform graph classification. Duvenaud et al. (2015) simply average all the node representations into a single vector. Gilmer et al. (2017) consider the node representations of the graph as a set and use the DeepSet aggregation (Zaheer et al. (2017)) to get a single representation. Li et al. (2015) add a virtual node to the graph which is connected to all the nodes and use the representation of the virtual node as the representation of the graph. Several approaches perform a deterministic hierarchical graph clustering step and combine the node representations in each cluster to learn hierarchical representations (Defferrard et al. (2016); Fey et al. (2018); Simonovsky and Komodakis (2017)). Instead of performing a deterministic clustering and then running a graph classification model, Ying et al. (2018b) learn the hierarchical structure jointly with the classifier in an end-to-end fashion.

3.2.2. DECODERS FOR LINK PREDICTION IN STATIC KGs

We provide an overview of the *translational*, *bilinear*, and *deep learning* decoders for KGs. When we discuss the expressivity of the decoders in this subsection, we assume the decoder is combined with a shallow encoder (see Section 3.1.2).

Translational decoders usually assume the encoder provides an embedding function such that $\text{EMB}(v) = (\mathbf{z}_v)$ for every $v \in \mathcal{V}$ where $\mathbf{z}_v \in \mathbb{R}^{d_1}$, and $\text{EMB}(r) = (\mathbf{z}_r, \mathbf{P}_r, \mathbf{Q}_r)$ for every $r \in \mathcal{R}$ where $\mathbf{z}_r \in \mathbb{R}^{d_2}$ and $\mathbf{P}_r, \mathbf{Q}_r \in \mathbb{R}^{d_1 \times d_2}$. That is, the embedding for a node contains a single vector whereas the embedding for a relation contains a vector and two matrices. For an edge (v, r, u) , these models use:

$$\|\mathbf{P}_r \mathbf{z}_v + \mathbf{z}_r - \mathbf{Q}_r \mathbf{z}_u\|_i \tag{18}$$




Figure 2: A graphical representation of the constraints over the P_r matrices for bilinear models (a) DistMult, (b) ComplEx, (c) CP, and (d) SimpleE taken from Kazemi and Poole (2018c) where lines represent the non-zero elements of the matrices. In ComplEx, the parameters represented by the dashed line are tied (i.e., equal) to the parameters represented by the solid line and the parameters represented by the dotted-and-dashed line are tied to the negative of the dotted line.

as the dissimilarity score for the edge where $\|\cdot\|_i$ represents norm i of a vector. i is usually either 1 or 2. Translational decoders differ in the restrictions they impose on \mathbf{P}_r and \mathbf{Q}_r . TransE (Bordes et al. (2013)) constrains $\mathbf{P}_r = \mathbf{Q}_r = \mathbf{I}_d$. So the dissimilarity function for TransE can be simplified to $\|\mathbf{z}_v + \mathbf{z}_r - \mathbf{z}_u\|_i$. In TransR (Lin et al. (2015)), $P_r = Q_r$. In STransE (Nguyen et al. (2016)), no restrictions are imposed on the matrices. Kazemi and Poole (2018c) proved that regardless of the encoder, TransE, TransR, STransE, and many other variants of translational approaches are not fully expressive for link prediction (see Definition 5 for a definition of fully expressive for link prediction) and identified severe restrictions on the type of relations these approaches can model.

Bilinear decoders usually assume the encoder provides an embedding function such that $\text{EMB}(\mathbf{v}) = (\mathbf{z}_v)$ for every $\mathbf{v} \in \mathcal{V}$ where $\mathbf{z}_v \in \mathbb{R}^d$, and $\text{EMB}(\mathbf{r}) = (\mathbf{P}_r)$ for every $\mathbf{r} \in \mathcal{R}$ where $\mathbf{P}_r \in \mathbb{R}^{d \times d}$. For an edge $(\mathbf{v}, \mathbf{r}, \mathbf{u})$, these models use:

$$\mathbf{z}'_v \mathbf{P}_r \mathbf{z}_u \quad (19)$$

as the similarity score for the edge. Bilinear decoders differ in the restrictions they impose on \mathbf{P}_r matrices (see Wang et al. (2018)). In RESCAL (Nickel et al. (2011)), no restrictions are imposed on the \mathbf{P}_r matrices. RESCAL is fully expressive with respect to link prediction, but the large number of parameters per relation makes RESCAL prone to overfitting. To reduce the number of parameters in RESCAL, DistMult (Yang et al. (2015)) constrains the \mathbf{P}_r matrices to be diagonal. This reduction in the number of parameters, however, comes at a cost: DistMult loses expressivity and is only able to model symmetric relations as it does not distinguish between the source and target vectors.

ComplEx (Trouillon et al. (2016)), CP (Hitchcock (1927)) and SimpleE (Kazemi and Poole (2018c)) reduce the number of parameters in RESCAL without sacrificing full expressivity. ComplEx extends DistMult by assuming the embeddings are complex (instead of real) valued, i.e. $\mathbf{z}_v \in \mathbb{C}^d$ and $\mathbf{P}_r \in \mathbb{C}^{d \times d}$ for every $\mathbf{v} \in \mathcal{V}$ and $\mathbf{r} \in \mathcal{R}$. Then, it slightly changes the score function to $\text{Real}(\mathbf{z}'_v \mathbf{P}_r \text{conjugate}(\mathbf{z}_u))$ where Real returns the real part of an imaginary number and conjugate takes an element-wise conjugate of the vector elements. By taking the conjugate of the target vector, ComplEx differentiates between source and target nodes

and does not suffer from the symmetry issue of DistMult. CP defines $\text{EMB}(v) = (\mathbf{z}_v, \bar{\mathbf{z}}_v)$, i.e. the embedding of a node consists of two vectors, where \mathbf{z}_v captures the v 's behaviour when it is the source of an edge and $\bar{\mathbf{z}}_v$ captures v 's behaviour when it is the target of an edge. For relations, CP defines $\text{EMB}(r) = (\mathbf{z}_r)$. The similarity function of CP for an edge (v, r, u) is then defined as $\mathbf{z}'_v \text{diag}(\mathbf{z}_r) \bar{\mathbf{z}}_u$. Realizing the information may not flow well between the two vectors of a node, SimpleE adds another vector to the relation embeddings as $\text{EMB}(r) = (\mathbf{z}_r, \bar{\mathbf{z}}_r)$ where $\bar{\mathbf{z}}_r$ models the behaviour of the inverse of the relation. Then, it changes the score function to be the average of $\mathbf{z}'_v \text{diag}(\mathbf{z}_r) \bar{\mathbf{z}}_u$ and $\mathbf{z}'_u \text{diag}(\bar{\mathbf{z}}_r) \bar{\mathbf{z}}_v$.

For ComplEx, CP, and SimpleE, it is possible to view the embedding for each node v as a single vector in \mathbb{R}^{2d} by concatenating the two vectors (for ComplEx, the two vectors correspond to the real and imaginary parts of the embedding vector). Then, the \mathbf{P}_r matrices can be viewed as being restricted according to Figure 2.

Other bilinear approaches include HolE (Sadilek and Kautz (2010)) whose equivalence to ComplEx has been established (see Hayashi and Shimbo (2017)), and Analogy (Liu et al. (2017)) where the \mathbf{P}_r matrices are constrained to be block-diagonal.

Deep learning-based decoders: Deep learning approaches typically use feed-forward or convolutional neural networks for scoring edges in a KG. Dong et al. (2014) and Santoro et al. (2017) consider $\text{EMB}(v) = (\mathbf{z}_v)$ for every node $v \in \mathcal{V}$ such that $\mathbf{z}_v \in \mathbb{R}^{d_1}$ and $\text{EMB}(r) = (\mathbf{z}_r)$ for every relation $r \in \mathcal{R}$ such that $\mathbf{z}_r \in \mathbb{R}^{d_2}$. Then for an edge (v, r, u) , they feed $[\mathbf{z}_v; \mathbf{z}_r; \mathbf{z}_u]$ (i.e., the concatenation of the three vector representations) into a feed-forward neural network that outputs a score for this edge. Dettmers et al. (2018) develop a score function based on convolutions. They consider $\text{EMB}(v) = (\mathbf{Z}_v)$ for each node $v \in \mathcal{V}$ such that $\mathbf{Z}_v \in \mathbb{R}^{d_1 \times d_2}$ and $\text{EMB}(r) = (\mathbf{Z}_r)$ for each relation $r \in \mathcal{R}$ such that $\mathbf{Z}_r \in \mathbb{R}^{d_1 \times d_2}$. For an edge (v, r, u) , first they combine \mathbf{Z}_v and \mathbf{Z}_r into a matrix $\mathbf{Z}_{vr} \in \mathbb{R}^{2d_1 \times d_2}$ by concatenating the two matrices on the rows, or by adding the i^{th} row of each matrix in turn. Then 2D convolutions with learnable filters are applied on \mathbf{Z}_{vr} generating multiple matrices and the matrices are vectorized into a vector \mathbf{c}_{vr} , where $|\mathbf{c}_{vr}|$ depends on the number of convolution filters. Then the score for the edge is computed as $(\mathbf{c}'_{vr} \mathbf{W}) \text{vec}(\mathbf{Z}_u)$ where $\mathbf{W} \in \mathbb{R}^{|\mathbf{c}_{vr}| \times (d_1 d_2)}$ is a weight matrix. Other deep learning approaches include HyperER (Balazevic et al. (2018)) which is another score function based on convolutions, and neural tensor networks (NTN) (Socher et al. (2013)) which contains feed-forward components as well as several bilinear components.

4. Encoders for Dynamic Graphs

In Section 3.1, we described different encoders for static graphs. In this section, we describe several general categories of encoders for dynamic graphs. Recall that reasoning problems for dynamic graphs can be for extrapolation or interpolation (see Section 2.3). Although some encoders may be used for both problems, the extrapolation and interpolation problems typically require different types of encoders. For extrapolation, one needs an encoder that provides node and relation embeddings based only on the observations in the past. For interpolation, however, at any time t , one needs an encoder that provides node and relation embeddings based on the observations before, at, and after t .

2. Alternatively, the matrices can be viewed as vectors of size $d_1 d_2$.

4.1. Aggregating Temporal Observations

A simple approach for dealing with the temporal aspect of a dynamic graph is through collapsing the dynamic graph into a static graph by aggregating the temporal observations (or the adjacency matrices) over time. Once an aggregated static graph is produced, a static encoder can be used to generate an embedding function.

Liben-Nowell and Kleinberg (2007) follow a simple aggregation approach for DTDGs by ignoring the timestamps and taking the sum (or union) of the entries of the adjacency matrices across all snapshots. That is, assuming $\mathbf{A}^1, \dots, \mathbf{A}^T$ represent the adjacency matrices for T timestamps, Liben-Nowell and Kleinberg (2007) first aggregate these adjacency matrices into a single matrix as follows:

$$\mathbf{A}_{sum}[i][j] = \sum_{t=1}^T \mathbf{A}^t[i][j] \quad (20)$$

Then a static decoder can be applied on \mathbf{A}_{sum} to learn an embedding function. Hisano (2018) also follows a similar aggregation scheme where he takes the union of the previous k formation and dissolution matrices of a DTDG. He defines the formation matrix for snapshot t as a matrix representing which edges have been added to the graph since $(t-1)^{\text{th}}$ snapshot and the dissolution matrix as a matrix representing which edges have been removed from the graph since $(t-1)^{\text{th}}$ snapshot. These simple approaches lose the timing information and may not perform well when timing information are of high importance.

An alternative to taking a uniform average of the adjacency matrices is to give more weights to snapshots that are more recent (Sharan and Neville (2008); Ibrahim and Chen (2015); Ahmed and Chen (2016); Ahmed et al. (2016)). Below is one such aggregation:

$$\mathbf{A}_{wsum}[i][j] = \sum_{t=1}^T \theta^{T-t} \mathbf{A}^t[i][j] \quad (21)$$

where $0 \leq \theta \leq 1$ controls the importance of recent snapshots. Larger values for θ put more emphasis on the more recent adjacency matrices.

Example 7 Let $\{\mathcal{G}^1, \mathcal{G}^2, \mathcal{G}^3\}$ be a DTDG with three snapshots. Let all \mathcal{G}^i s have the same set $\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$ of nodes and the adjacency matrices be as follows:

$$\mathbf{A}^1 = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} \quad \mathbf{A}^2 = \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix} \quad \mathbf{A}^3 = \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

The aggregation scheme in Equation (20) and Equation (21) (assuming $\theta = 0.5$) respectively aggregate the three adjacency matrices into \mathbf{A}_{sum} and \mathbf{A}_{wsum} as follows:

$$\mathbf{A}_{sum} = \begin{bmatrix} 0 & 3 & 2 \\ 3 & 0 & 2 \\ 2 & 2 & 0 \end{bmatrix} \quad \mathbf{A}_{wsum} = \begin{bmatrix} 0 & \frac{7}{4} & \frac{3}{2} \\ \frac{7}{4} & 0 & \frac{3}{4} \\ \frac{3}{2} & \frac{3}{4} & 0 \end{bmatrix}$$

Then an embedding function can be learned using \mathbf{A}_{sum} or \mathbf{A}_{wsum} (e.g., by using decomposition approaches). Although the interaction evolution between \mathbf{v}_1 and \mathbf{v}_3 (which were not

connected at the beginning, but then formed a connection) is different from the interaction evolution between v_2 and v_3 (which were connected at the beginning and then got disconnected), \mathbf{A}_{sum} assigns the same number to both these pairs. \mathbf{A}_{wsum} contains more temporal information compared to \mathbf{A}_{sum} , but still loses large amounts of information. For instance, it is not possible to realize from \mathbf{A}_{wsum} that v_2 and v_3 got disconnected only recently.

The approaches based on aggregating temporal observations typically enjoy advantages such as simplicity, scalability, and the capability to directly use a large body of literature on learning from static graphs. The aggregation in Equation (20) can be potentially used for both interpolation and extrapolation. The aggregation in Equation (20) may be more suited for extrapolation as it weighs recent snapshots more than the old ones. However, it can be easily adapted for interpolation by, e.g., changing θ^{T-t} in the equation with $\theta^{|t_q-t|}$ where t_q is the timestamp for which we wish to make a prediction and $|\cdot|$ returns the absolute value. Note that the aggregation approaches may lose large amounts of useful information hindering them from making accurate predictions in many scenarios.

4.2. Aggregating Static Features

Rather than first aggregating a dynamic graph over time to produce a static graph and then running static encoders on the aggregated graph, in the case of DTDGs, one may first apply a static encoder to each snapshot and then aggregate the results over time. Let $\{\mathcal{G}^1, \dots, \mathcal{G}^T\}$ be a DTDG. The main idea behind the approaches in this category is to first use a static encoder (e.g., an encoder from Section 3.1.1) to compute/learn node features \mathbf{z}_v^t for each node v at each timestamp t . The features for each timestamp are computed/learned independently of the other timestamps. Then, these features are aggregated into a single feature vector that can be fed into a decoder.

Yao et al. (2016) aggregate features into a single feature vector as follows:

$$\mathbf{z}_v = \sum_{t=1}^T \exp(-\theta(T-t)) \mathbf{z}_v^t \quad (22)$$

thus exponentially decaying older features. Zhu et al. (2012) follow a similar strategy where they compute features for each pair of nodes and take a weighted sum (with prefixed weights) of the features, giving higher weights to the features coming from more recent snapshots.

Rather than using an explicitly defined aggregator (e.g., exponential decay) that assigns prefixed weights to previous snapshots, one can fit a time-series model to the features from previous snapshots and use this model to predict the values of the features for the next snapshot. For the time-series model, Huang and Lin (2009) and Güneş et al. (2016) use the ARIMA model (Box et al. (2015)), da Silva Soares and Prudêncio (2012) use ARIMA and other models such as moving averages, and Moradabadi and Meybodi (2017) use an approach based on some basic reinforcement learning.

Example 8 Consider the DTDG in Example 7. A simple example of creating node embeddings by aggregating static features is to use the common neighbor static encoder (see Section 3.1.1 for details) to obtain node embeddings at each timestamp and then combine

these embeddings using Equation (22). The common neighbors encoder applied to $\mathbf{A}^1, \mathbf{A}^2$ and \mathbf{A}^3 gives the following embeddings:

$$\mathbf{Z}^1 = \mathbf{S}_{CN}^1 = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 2 & 0 \\ 1 & 0 & 1 \end{bmatrix} \quad \mathbf{Z}^2 = \mathbf{S}_{CN}^2 = \begin{bmatrix} 2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{bmatrix} \quad \mathbf{Z}^3 = \mathbf{S}_{CN}^3 = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 1 & 1 \end{bmatrix}$$

In the above matrices, $\mathbf{Z}^t[i]$ corresponds to the embedding of \mathbf{v}_i at timestamp t . Notice how the embeddings for different timestamps are computed independently of the other timestamps. Also note that while we used the common neighbors encoder, any other static encoder from Section 3.1 could potentially be used. Considering a value of $\theta = 0.5$, the above embeddings are then aggregated into:

$$\mathbf{Z}^{aggr} \approx \begin{bmatrix} 3.58 & 0.61 & 0.97 \\ 0.61 & 2.95 & 1.61 \\ 0.97 & 1.61 & 2.58 \end{bmatrix}$$

where $\mathbf{Z}^{aggr}[i]$ corresponds to the embedding for \mathbf{v}_i .

Scalability: Depending on the number of snapshots and the static encoder used for feature generation, the approaches that compute node features/embeddings at each snapshot independently of the other snapshots and then aggregate these features may be computationally expensive. In the upcoming subsections, for some choices of static encoders (e.g., for decomposition and random-walk approaches), we will see some techniques to save computations in later snapshots by leveraging the computations from the previous snapshots.

4.3. Time as a Regularizer

A common approach to leverage the temporal aspect of DTDGs is to use time as a regularizer to impose a smoothness constraint on the embeddings of each node over consecutive snapshots (Chakrabarti et al. (2006); Chi et al. (2009); Kim and Han (2009); Gupta et al. (2011); Yao et al. (2016); Zhu et al. (2016); Zhou et al. (2018b)). Consider a DTDG as $\{\mathcal{G}^1, \dots, \mathcal{G}^T\}$. For a node \mathbf{v} , let $\text{EMB}^{t-1}(\mathbf{v}) = (\mathbf{z}_\mathbf{v}^{t-1})$ represent the vector representation learned for this node at the $(t-1)^{\text{th}}$ snapshot. To learn the vector representation for \mathbf{v} at the t^{th} snapshot, the approaches in this class typically use a static encoder to learn an embedding function for \mathcal{G}^t with the additional constraint that for any node $\mathbf{v} \in \mathcal{V}^t$ such that $\mathbf{v} \in \mathcal{V}^{t-1}$ (i.e. for any node \mathbf{v} that has been in the graph in the previous and current snapshots), $\text{dist}(\mathbf{z}_\mathbf{v}^{t-1}, \mathbf{z}_\mathbf{v}^t)$ should be small, where dist is a distance function. This constraint is often called the *smoothness* constraint. A common choice for the distance function is the Euclidean distance:

$$\text{dist}(\mathbf{z}_\mathbf{v}^t, \mathbf{z}_\mathbf{v}^{t-1}) = \|\mathbf{z}_\mathbf{v}^t - \mathbf{z}_\mathbf{v}^{t-1}\| \quad (23)$$

but distance in other spaces may also be used (see, e.g., Chi et al. (2009)). Singer et al. (2019) add a rotation projection to align the embedding $\mathbf{z}_\mathbf{v}^t$ s with the embedding $\mathbf{z}_\mathbf{v}^{t-1}$ s before taking the Euclidean distance. Their distance function can be represented as follows:

$$\text{dist}(\mathbf{z}_\mathbf{v}^t, \mathbf{z}_\mathbf{v}^{t-1}) = \|\mathbf{R}^t \mathbf{z}_\mathbf{v}^t - \mathbf{z}_\mathbf{v}^{t-1}\| \quad (24)$$

where \mathbf{R}^t is a rotation matrix. Instead of Euclidean distance, Milani Fard et al. (2019) define the `dist` function based on the angle between the two vectors. Their distance function can be written as follows:

$$\text{dist}(\mathbf{z}_v^t, \mathbf{z}_v^{t-1}) = 1 - (\mathbf{z}_v^t)' \mathbf{z}_v^{t-1} \quad (25)$$

where all embedding vectors are restricted to have a norm of 1. Note that the smaller the angle between \mathbf{z}_v^t and \mathbf{z}_v^{t-1} , the closer $(\mathbf{z}_v^t)' \mathbf{z}_v^{t-1}$ is to 1 and so the closer $\text{dist}(\mathbf{z}_v^t, \mathbf{z}_v^{t-1}) = 1 - (\mathbf{z}_v^t)' \mathbf{z}_v^{t-1}$ is to 0. Liu et al. (2019) also use time as a regularizer, but they turn the representation learning problem into a constrained optimization problem that can be approximated in a reasonable amount of time. As new observations are made, their representations can be updated in a short amount of time and so their model may be used for streaming scenarios. The model they propose also handles addition of new nodes to the graph. Pei et al. (2016) propose a dynamic factor graph model for node classification in which they use the temporal information in a similar way as the other approaches in this section: they impose factors that decrease the probability of the worlds where the label of a node at the t^{th} snapshot is different from the previous snapshots (exponentially decaying the importance of the labels for the older snapshots). Using time as a regularizer can be useful for both interpolation and extrapolation problems.

Example 9 Consider the DTDG in Example 7. Suppose we want to provide node embeddings by using a static autoencoder approach (see Section 3.1.5 for details) while using time as a regularizer. In the first timestamp, we train an autoencoder whose encoder takes as input $\mathbf{A}^1[i]$, feeds it through its encoder and generates $\mathbf{z}_{v_i}^1$, and then feeds $\mathbf{z}_{v_i}^1$ through its reconstructor to output $\hat{\mathbf{A}}^1[i]$ with the loss function being $\sum_{i=1}^{|\mathcal{V}|} \|\mathbf{A}^1[i] - \hat{\mathbf{A}}^1[i]\|$. In the second timestamp, we follow a similar approach but instead of the loss function being $\sum_{i=1}^{|\mathcal{V}|} \|\mathbf{A}^2[i] - \hat{\mathbf{A}}^2[i]\|$, we define the loss function to be $\sum_{i=1}^{|\mathcal{V}|} \|\mathbf{A}^2[i] - \hat{\mathbf{A}}^2[i]\| + \|\mathbf{z}_{v_i}^2 - \mathbf{z}_{v_i}^1\|$. We continue a similar procedure in the third timestamp by defining the loss function as $\sum_{i=1}^{|\mathcal{V}|} \|\mathbf{A}^3[i] - \hat{\mathbf{A}}^3[i]\| + \|\mathbf{z}_{v_i}^3 - \mathbf{z}_{v_i}^2\|$. Note that here we are using the distance function from Equation (23) but other distance function can be used as well.

Imposing smoothness constraints through penalizing the distance between the vector representations of a node at consecutive snapshots stops the vector representation from having sharp changes. While this may be desired for some applications, in some other applications a node may change substantially from one snapshot to the other. As an example, if a company gets acquired by a large company, it is expected that its vector representation in the next snapshot makes sharp changes. Instead of penalizing the distance of the vector representations for a node at consecutive snapshots, one may simply initialize the representations (or the model) for time t with the learned representations (or model) at time $t - 1$ and then allow the static encoder to further optimize the representation at time t (see, e.g., Goyal et al. (2017)). This procedure implicitly imposes the smoothness constraint while also allowing for sharp changes when necessary.

Another notable work where time is used as a regularizer is an extension of a well-known model for static graphs, named *LINE* (Tang et al. (2015)), to DTDGs by Du et al. (2018). Besides using time as regularizer, the authors propose a way of recomputing the node embeddings only for the nodes that have been influenced greatly from the last snapshot.

4.4. Decomposition-based Encoders

A good application of decomposition methods to dynamic graphs is to use them as an alternative to aggregating temporal observations described in Section 4.1. Consider a DTDG as $\{\mathcal{G}^1, \dots, \mathcal{G}^T\}$ such that $\mathcal{V}^1 = \mathcal{V}^2 = \dots = \mathcal{V}^T = \mathcal{V}$ (i.e. nodes are not added or removed). As was proposed by Dunlavy et al. (2011), the adjacency matrices $\mathbf{A}^1, \dots, \mathbf{A}^T$ for T timestamps can be stacked into an order 3 tensor $\mathcal{A} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}| \times T}$. Then one can do a d -component tensor decomposition (e.g., CP decomposition, see Harshman et al. (1970)):

$$\mathcal{A} \approx \sum_{k=1}^d \lambda_k \mathbf{a}_k \otimes \mathbf{b}_k \otimes \mathbf{c}_k \quad (26)$$

where $\lambda_k \in \mathbb{R}_+$, $\mathbf{a}_k, \mathbf{b}_k \in \mathbb{R}^{|\mathcal{V}|}$, $\mathbf{c}_k \in \mathbb{R}^T$, and \otimes is a tensor product of vector spaces. The temporal pattern is captured in the \mathbf{c}_k s, and a combination of \mathbf{a}_k s and \mathbf{b}_k s can be used as the node (or edge) embeddings. These embeddings can be used to make predictions about any time $1 \leq t \leq T$, i.e. for interpolation. For extrapolation, Dunlavy et al. (2011) used the Holt-Winters method (see, Chatfield and Yar (1988)): given the input \mathbf{c}_k , it predicts an L -dimensional vector \mathbf{c}'_k , which is the prediction of the temporal factor for the next L timesteps. Then they predict the adjacency tensor for the next L snapshots as $\hat{\mathcal{A}} = \sum_{k=1}^d \lambda_k \mathbf{a}_k \otimes \mathbf{b}_k \otimes \mathbf{c}'_k$. One can also use other forms of tensor decomposition, e.g. Tucker decomposition or HOSVD (Rabanser et al. (2017)). Xiong et al. (2010) propose a probabilistic factorization of \mathcal{A} where the nodes are represented as normal distributions with the means coming from \mathbf{a}_k s and \mathbf{b}_k s. They also impose a smoothness prior over the temporal vectors corresponding to using time as a regularizer (see Section 4.3). After some time steps, one needs to update the tensor decomposition for more accurate future predictions. The recomputation can be quite costly so one can try incremental updates (see Gujral et al. (2018) and Letourneau et al. (2018)).

Yu et al. (2017a) present another way of incorporating temporal dependencies into the embeddings with decomposition methods. As above, let $\mathbf{A}^1, \dots, \mathbf{A}^T$ be the adjacency matrices for T timestamps. Yu et al. (2017a) predict $\hat{\mathbf{A}}^{T+l}$, where $l \in \mathbb{N}$, as follows. First, they solve the optimization problem:

$$\min \sum_{t=T-\omega}^T e^{-\theta(T-t)} \|\mathbf{A}^t - \mathbf{U}(\mathbf{V}^t)'(\mathbf{P}^t)'\|_F^2, \quad (27)$$

where $\mathbf{P}^t = (\mathbf{I} - \alpha)(\mathbf{I} - \alpha\sqrt{\mathbf{D}^t}\mathbf{A}^t\sqrt{\mathbf{D}^t})^{-1}$ is the projection onto feature space that ensures neighboring nodes have similar feature vectors (see Yu et al. (2017a) for details), ω is a window of timestamps into consideration, $\alpha \in (0, 1)$ is a regularization parameter, θ is a decay parameter, $\mathbf{U} \in \mathbb{R}^{|\mathcal{V}| \times d}$ is a matrix that does not depend on time, and $\mathbf{V}^t \in \mathbb{R}^{|\mathcal{V}| \times d}$ is a matrix with explicit time dependency (in the paper, it is a polynomial in time with matrix coefficients). The optimization problem can be slightly rewritten using the sparsity of \mathbf{A} and then solved with stochastic gradient descent. The prediction can be obtained as $\hat{\mathbf{A}}^{T+l} = \mathbf{U}(\mathbf{V}^{T+l})'$. From the point of view of the encoder-decoder framework, \mathbf{U} can be interpreted as static node features and \mathbf{V}^t s are time-dependent features (one takes the i^{th} row of the matrices as an embedding of the i^{th} node).

One can extend the tensor decomposition idea to the case of temporal KGs by modelling the KG as an order 4 tensor $\mathcal{A} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{R}| \times |\mathcal{V}| \times T}$ and decomposing it using CP, Tucker, or some other decomposition approach to obtain entity, relation, and timestamp embeddings (see Tresp et al. (2015); Esteban et al. (2016); Tresp et al. (2017)). Tresp et al. (2015) and Tresp et al. (2017) study the connection between these order 4 tensor decomposition approaches and the human cognitive memory functions.

The streaming scenario. As was discussed in Subsection 3.1.3, one can learn node embedding using either eigen-decomposition or SVD for graph matrices for each timestamp. Then one can aggregate these features as in Section 4.2 for predictions. However, recalculating decomposition at every timestamp may be computationally expensive. So one needs to come up with incremental algorithms that will update the current state in the streaming case.

Incremental eigenvalue decomposition (Chen and Tong (2015); Li et al. (2017a); Wang et al. (2017b)) is based on perturbation theory. Consider a generalized eigenvalue problem as in Equation (9). Then assume that in the next snapshot we add a few new edges to the graph \mathcal{G}^T . In this case, the Laplacian and the degree matrix change by a small amount: $\Delta \mathbf{L}$ and $\Delta \mathbf{D}$ respectively. Assume that we have solved Equation (9) and $\{(\lambda_i, \mathbf{y}_i)\}_{i=1}^{|\mathcal{V}|}$ is the solution. Then one can find the solution to the new generalized eigenvalue problem for the graph \mathcal{G}^{T+1} in the form: updated eigenvalues $\approx \lambda + \Delta \lambda$ and updated eigenvectors $\approx \mathbf{y} + \Delta \mathbf{y}$, where $\Delta \lambda$ and $\Delta \mathbf{y}$ can be efficiently computed. For example,

$$\Delta \lambda_i = \frac{\mathbf{y}_i' \Delta \mathbf{L} \mathbf{y}_i - \lambda_i \mathbf{y}_i' \Delta \mathbf{D} \mathbf{y}_i}{\mathbf{y}_i' \mathbf{D} \mathbf{y}_i}. \quad (28)$$

An analogous formula could be written for $\Delta \mathbf{y}_i$. The Davis-Kahan theorem (Davis and Kahan (1970)) gives an approximation error for the top d eigen-pairs.

As shown by Levin et al. (2018), one can recalculate the adjacency spectral embedding (see Section 3.1.3 for the construction) in case of addition of a new node v to a graph \mathcal{G} . Denote $\mathbf{a}_v \in \mathbb{R}^{|\mathcal{V}|}$ a binary vector where each entry indicates whether there is an edge between the added node and an already existing node. One can find \mathbf{z}_v as the solution to the maximum likelihood problem to fit $\mathbf{a}_v \sim \text{Bernoulli}(\mathbf{Z}' \mathbf{z}_v)$, where \mathbf{Z} is as in Formula (10).

Brand (2006) proposes an efficient way to update the singular value decomposition of a matrix \mathbf{S} when another lower rank matrix of the same size $\Delta \mathbf{S}$ is added to it. Consider the problem in Equation (12). If one knows the solution $(\mathbf{U}_s, \mathbf{U}_t, \mathbf{\Sigma})$ and $\Delta \mathbf{S}$ is an update of the matrix, one can find a general formula for the update of the SVD using some basic computations with block matrices. However, this becomes especially efficient if we approximate the increment as a rank one matrix: $\Delta \mathbf{S} = \mathbf{a} \mathbf{b}'$ (see also Stange (2008)). Bunch and Nielsen (1978) studied how SVD can be updated when a row or column is added to or removed from a matrix \mathbf{S} . This can be applied to get the encoding for a DTDG in the case of node addition or deletion.

One problem with incremental updates is that the approximation error keeps accumulating gradually. As a solution, one needs to recalculate the model from time to time. However, since the recalculation is expensive, one needs to find a proper time when the error becomes intolerable. Usually in applications people use heuristic methods (e.g., restart after a certain time); however, ideally a timing should depend on the graph dynamics. Zhang et al. (2018d) propose a new method where given a tolerance threshold, it notifies at what timestamp the approximation error exceeds the threshold.

4.5. Random Walk Encoders

Recently, several approaches have been proposed to leverage or extend the random walk models for static graphs to dynamic graphs. In this section, we provide an overview of these approaches.

Consider a DTDG as $\{\mathcal{G}^1, \dots, \mathcal{G}^T\}$. Mahdavi et al. (2018) first generate random walks on \mathcal{G}^1 similar to the random walk models on static graphs and then feed those random walks to a model \mathcal{M}^1 that learns to produce vector representations for nodes given the random walks. For the t^{th} snapshot ($t > 1$), instead of generating random walks from scratch, they keep the *valid* random walks from $(t - 1)^{\text{th}}$ snapshot, where they define a random walk as valid if all its nodes and the edges taken along the walk are still in the graph in the t^{th} snapshot. They generate new random walks only starting from the affected nodes, where affected nodes are the nodes that have been either added in this snapshot, or are involved in one or more edge addition or deletion. Having obtained the updated random walks, they initialize \mathcal{M}^t with the learned parameters from \mathcal{M}^{t-1} and then allow \mathcal{M}^t to be optimized and produce the node embeddings for the t^{th} snapshot.

Bian et al. (2019) take a strategy similar to that of Mahdavi et al. (2018) but for KGs. They use metapath2vec (explained in Section 3.1.4) to generate random walks on the initial KG. Then, at each snapshot, they use metapath2vec to generate random walks for the affected nodes and re-compute the embeddings for these nodes.

Sajjad et al. (2019) observed that by keeping the valid random walks from the previous snapshot and naively generating new random walks starting from the affected nodes, the resulting random walks may be biased. That is, the random walks obtained by following this procedure may have a different distribution than generating random walks for the new snapshot from scratch. Example 10 demonstrates one such example.

Example 10 Consider Figure 1(a) as the first snapshot of a DTDG and assume the following random walks have been generated for this graph (two random walks starting from each node) following a uniform transition:

- | | | |
|--------------------|--------------------|--------------------|
| 1) v_1, v_2, v_1 | 2) v_1, v_2, v_3 | 3) v_2, v_1, v_3 |
| 4) v_2, v_3, v_1 | 5) v_3, v_2, v_1 | 6) v_3, v_1, v_2 |

Now assume the graph in Figure 1(b) represents the next snapshot. The affected nodes are v_2 , which has a new edge, and v_4 , which has been added in this snapshot. A naive approach for updating the above set of random walks is to remove random walks 3 and 4 (since they start from an affected node) and add two new random walks from v_2 and two from v_4 . This may give the following eight walks:

- | | | | |
|--------------------|--------------------|--------------------|--------------------|
| 1) v_1, v_2, v_1 | 2) v_1, v_2, v_3 | 3) v_2, v_4, v_2 | 4) v_2, v_3, v_1 |
| 5) v_3, v_2, v_1 | 6) v_3, v_1, v_2 | 7) v_4, v_2, v_3 | 8) v_4, v_2, v_1 |

In the above 8 random walks, the number of times a transition from v_2 to v_4 has been made is 1 and the number of times a transition from v_2 to v_1 (or v_3) has been made is 3, whereas, if new random walks are generated from scratch, the two numbers are expected to be the same. The reason for this bias is that in random walks 1, 2, 5, and 6, the walk could not go from v_2 to v_4 as v_4 did not exist when these walks were generated. Note that performing more random walks from each node does not solve the bias problem.

Sajjad et al. (2019) propose an algorithm for generating unbiased random walks for a new snapshot while reusing the valid random walks from the previous snapshot. NetWalk (Yu et al. (2018b)) follows a similar approach as the previous two approaches. However, rather than relying on natural language processing techniques to generate vector representations for nodes given random walks, they develop a customized autoencoder model that learns the vector representations for nodes while minimizing the pairwise distance among the nodes in each random walk.

The previous three approaches mainly leverage the temporal aspect of DTDGs to reduce the computations. They can be useful in the case of feature aggregation (see Section 4.2) when random walk encoders are used to learn features at each snapshot. However, they may fail at capturing the evolution and the temporal patterns of the nodes. Nguyen et al. (2018b,a) propose an extension of the random walk models for CTDGs that also captures the temporal patterns of the nodes.

Consider a CTDG as $(\mathcal{G}, \mathcal{O})$ where the only type of event in \mathcal{O} is the addition of new edges. Therefore, the nodes are fixed and each element of \mathcal{O} can be represented as $(AddEdge, (v, u), t_{(v,u)})$ indicating an edge was added between v and u at time $t_{(v,u)}$. Nguyen et al. (2018b,a) constrain the random walks to respect time, where they define a random walk on a CTDG that respects time as a sequence v_1, v_2, \dots, v_l of nodes where:

$$v_i \in \mathcal{V} \text{ for all } 1 \leq i \leq l \tag{29}$$

$$(AddEdge, (v_i, v_{i+1}), t_{(v_i, v_{i+1})}) \in \mathcal{O} \text{ for all } 1 \leq i \leq l - 1 \tag{30}$$

$$t_{(v_i, v_{i+1})} \leq t_{(v_{i+1}, v_{i+2})} \text{ for all } 1 \leq i \leq l - 2 \tag{31}$$

That is, the sequence of edges taken by each random walk only moves forward in time. Similar to the random walks on static graphs, the initial node to start a random walk from and the next node to transition to can come from a distribution. Unlike the static graphs, however, these distributions can be a function of time. For instance, consider a walk that has currently reached a node u by taking an edge (v, u) that has been added at time t . The edge for the next transition (to be selected from the outgoing edges of u that have been added after t) can be selected with a probability proportional to how long after t they were added to the graph.

Example 11 *Assume $t_1 < t_2 < t_3 < t_4 < t_5$ for the CTDG in Figure 1(d). Consider a random walk that has started at v_1 , then transitioned to v_2 , and is now deciding the next node to transition to. According to Nguyen et al. (2018b)'s strategy, even though both v_4 and v_5 are neighbors of v_2 , only the transition to v_4 is valid as the edge between v_2 and v_4 has been added after the edge between v_1 and v_2 whereas the edge between v_2 and v_5 has been added before the edge between v_1 and v_2 .*

De Winter et al. (2018) follow a similar approach as Nguyen et al. (2018b) but for DTDGs instead of CTDGs. Their experiments show that in some cases, discretizing a CTDG into a DTDG and then running the random walks that respect time on the DTDG results in better performance. Bastas et al. (2019) also follow a similar approach as Nguyen et al. (2018b), but they divide the time horizon into two intervals one corresponding to observations before some time t_s and the other corresponding to the observations after t_s . They aggregate the observations until time t_s into a static graph (see Section 4.1) following the intuition that

older observations mainly contain topological information and not temporal information. They run static random walks on the static graph from the first interval and temporal random walks that respect time (with custom distributions for selecting the initial node and the node to transition to) on the second interval. Both the static and temporal walks are then used to learn node embeddings.

4.5.1. ANALYSIS OF RANDOM WALK ENCODERS

Supervised and unsupervised learning: Similar to decomposition and autoencoder-based approaches, one of the major advantages of the encoders based on random walks is that they provide an embedding function without needing to be combined with a decoder. Therefore, the encoder can be used for unsupervised learning approaches such as clustering and community detection (Xin et al. (2016); Yu et al. (2017b)). However, the disconnect between the encoder and the decoder typically prevents these models from being trained end-to-end. Therefore, for supervised prediction tasks, the embedding learned for nodes are not optimized for the prediction problem.

The streaming scenario: When new observations are made, random walk approaches typically require to perform new walks that take the new observations into account and then update the node embeddings based on the new walks. This update usually requires a few rounds of computing gradients, which, depending on the size of the dynamic graph, can be quite time-consuming. This makes random walk approaches not an ideal option for the streaming scenario.

Random walk for attributed graphs: Using random walks for representation learning has been mostly done for non-attributed graphs. An interesting direction for future research is to extend the approaches discussed in this section to attributed graphs.

4.6. Sequence-Model Encoders

A natural choice for modeling dynamic graphs is by extending sequence models to graph data. With the success of RNNs in several synchronous sequence modeling problems (Mikolov et al. (2010); Bahdanau et al. (2014); Hermann et al. (2015); Mesnil et al. (2015); Huang et al. (2015b); Seo et al. (2016)), where the duration between any two consecutive items in the sequence is considered equal, and several asynchronous sequence modeling problems (Choi et al. (2016); Li et al. (2017c); Du et al. (2016); Neil et al. (2016); Zhu et al. (2017); Hu and Qi (2017)), RNNs have been a common choice for extending sequence models to DTDGs and CTDGs. In the next subsections, we describe the RNN-based models for DTDGs, which can be considered a synchronous sequence modeling problem, and CTDGs, which can be considered an asynchronous sequence modeling problem. We also describe some other sequence modeling approaches that have been extended to dynamic graphs. Note that sequence model approaches are mainly designed for extrapolation as they go through the observations sequentially and provide embeddings at the current time based on the past. However, one may use them for interpolation by, e.g., using bi-directional sequence models one running forward and providing embeddings based on everything before some time t and the other running backward and providing embeddings based on everything after time t .

4.6.1. RNN-BASED ENCODERS FOR DTDGS

Consider a DTDG as $\{\mathcal{G}^1, \mathcal{G}^2, \dots, \mathcal{G}^T\}$. Let \mathcal{M} be a (differentiable) encoder, which, given a static graph \mathcal{G}^t , outputs a vector representation for each node. As an example, \mathcal{M} can be a GCN (see Section 3.1.6 for details).

One way of leveraging RNNs for DTDGs is as follows. We run \mathcal{M} on each \mathcal{G}^t and obtain a sequence $\mathbf{z}_v^1, \mathbf{z}_v^2, \dots, \mathbf{z}_v^T$ of vector representations for each node v . This sequence is then fed into an RNN that produces a vector representation \mathbf{z}_v for v containing information from v 's history and evolution. These vector representations of the nodes can then be fed into a decoder to make predictions about the nodes. The idea behind this approach is similar to the static feature aggregation idea described in Section 4.2 except that the weights of the RNN and the model \mathcal{M} are learned simultaneously and over all the snapshots. The t^{th} step of the encoder for this architecture can be represented as:

$$\mathbf{z}_{v_1}^t, \dots, \mathbf{z}_{v_{|\mathcal{V}^t|}}^t = \mathcal{M}(\mathcal{G}^t) \tag{32}$$

$$\mathbf{h}_{v_j}^t = \text{RNN}(\mathbf{h}_{v_j}^{t-1}, \mathbf{z}_{v_j}^t) \text{ for } j \in [1, |\mathcal{V}^t|] \tag{33}$$

which can be equivalently represented as:

$$\mathbf{Z}^t = \mathcal{M}(\mathcal{G}^t) \tag{34}$$

$$\mathbf{H}^t = \text{RNN}(\mathbf{H}^{t-1}, \mathbf{Z}^t) \tag{35}$$

where $\mathbf{Z}^t \in \mathbb{R}^{|\mathcal{V}^t| \times d}$ represents the vector representations of size d for the $|\mathcal{V}^t|$ nodes in the graph at snapshot t and $\mathbf{H}^t \in \mathbb{R}^{|\mathcal{V}^t| \times d}$ represents the hidden state of the RNN corresponding to vector representations of size d for the $|\mathcal{V}^t|$ nodes in the graph that captures the history of the nodes as well. In this architecture, \mathcal{M} aims at capturing the structural information for each node at each snapshot and the RNN aims at capturing the temporal information. The approach described above has been proposed and used in different works. Model 1 of Seo et al. (2018) uses this approach where \mathcal{M} is the GCN proposed by Defferrard et al. (2016) and the RNN is a standard LSTM. Narayan and Roe (2018) also use this approach with \mathcal{M} being the GCN proposed by Niepert et al. (2016) and the RNN being a standard LSTM. Manessi et al. (2017) modify this approach slightly by (mainly) adding skip-connections in the GCN part. Another similar architecture is proposed by Mohanty and Pozdnukhov (2018). Instead of obtaining $\mathbf{z}_{v_j}^t$ s by running a GCN that aggregates the features of neighboring nodes only at the t^{th} snapshot, Yu et al. (2019) propose a 3D GCN that aggregates the features of neighboring nodes on a window of previous snapshots (i.e., the aggregation is both spatial and temporal rather than just being spatial).

In the above approach, \mathcal{M} is independent of the RNN. That is, the vector representations for nodes provided by \mathcal{M} are independent of the node histories captured in $\mathbf{h}_{v_j}^t$ s. The *embedded approaches* aim at embedding the model(s) \mathcal{M} into the RNN so that \mathcal{M} can also use the node histories.

One such embedded approach has been proposed by Chen et al. (2018a), where the authors combine the GCN proposed by Defferrard et al. (2016) with LSTMs. Consider a DTDG as $\{\mathcal{G}^1, \mathcal{G}^2, \dots, \mathcal{G}^T\}$, let \mathbf{A}^t represent the adjacency matrix for \mathcal{G}^t , and let $\mathbf{A}^t[j]$ represent the j^{th} row of \mathbf{A}^t corresponding to the neighborhood of node v_j . Let $\mathbf{C}^{t-1}[j]$ and $\mathbf{H}^{t-1}[j]$ represent the memory and hidden state of the LSTM at time $t - 1$ for node v_j . Let

$\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3, \mathcal{M}_4$ and \mathcal{M}_5 be five GCN models (same model with different parameters), where the node representations for $\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_4$ and \mathcal{M}_5 are initialized according to \mathbf{H}^{t-1} and for \mathcal{M}_3 are initialized according to \mathbf{C}^{t-1} . Let $\mathcal{M}_i(\mathcal{G})[j]$ represent the vector representation provided by \mathcal{M}_i for node \mathbf{v}_j when applied to \mathcal{G} . The embedded model of Chen et al. (2018a) can be formulated as:

$$\mathbf{i}^t = \sigma(\mathbf{W}_{ii}\mathbf{A}^t[j] + \mathcal{M}_1(\mathcal{G}^t)[j] + \mathbf{b}_i) \quad (36)$$

$$\mathbf{f}^t = \sigma(\mathbf{W}_{fi}\mathbf{A}^t[j] + \mathcal{M}_2(\mathcal{G}^t)[j] + \mathbf{b}_f) \quad (37)$$

$$\mathbf{C}^t[j] = \mathbf{f}_t \odot \mathcal{M}_3(\mathcal{G}^t)[j] + \mathbf{i}^t \odot \text{Tanh}(\mathbf{W}_{ci}\mathbf{A}^t[j] + \mathcal{M}_4(\mathcal{G}^t)[j] + \mathbf{b}_c) \quad (38)$$

$$\mathbf{o}^t = \sigma(\mathbf{W}_{oi}\mathbf{A}^t[j] + \mathcal{M}_5(\mathcal{G}^t)[j] + \mathbf{b}_o) \quad (39)$$

$$\mathbf{H}^t[j] = \mathbf{o}^t \odot \text{Tanh}(\mathbf{C}^t[j]) \quad (40)$$

where the above formulation is done for all $1 \leq j \leq |\mathcal{V}^t|$. The \mathcal{M}_i models are embedded into the LSTM gates and use its memory and hidden state in their computations. The above formulation can be considered as a standard LSTM taking a sequence of adjacency matrices as input, where the gate computations for the memory and hidden states have been replaced with GCN operations to take the graph structure into account. Other similar ways of embedding the \mathcal{M} models into RNNs can be seen in Li et al. (2017b) and Seo et al. (2018). Instead of embedding a GCN into an RNN, Pareja et al. (2019) embed an RNN into a GCN by running a GCN with different parameters at each snapshot where the RNN provides the weights of the GCN at the t^{th} snapshot based on the weights of the GCNs at the previous snapshots.

4.6.2. OTHER SEQUENCE-MODEL ENCODERS FOR DTDGS

Besides RNNs, other sequence models have been also used as encoders for DTDGs. Sarkar et al. (2007) use Kalman filters to track the embeddings of the nodes through time for a bipartite graph. Each timestep of the Kalman filter corresponds to a snapshot of the DTDG. The t^{th} timestep observes the adjacency matrix \mathbf{A}^t corresponding to the t^{th} snapshot \mathcal{G}^t of the DTDG and updates the node embeddings accordingly. The observations for each element \mathbf{A}^t are considered to be independent of each other and $\mathbf{A}^t[i][j]$ is defined to be proportional to the distance between the embedding vectors of \mathbf{v}_i and \mathbf{v}_j . To make the computations tractable, this conditional probability of the observation (i.e., the adjacency matrix) given the hidden state (i.e., the embeddings) is approximated by a Gaussian.

Inspired by fully attentive models (see Section 2.6), Sankar et al. (2018) propose DySAT: a fully attentive model for DTDGs. Consider a DTDG as $\{\mathcal{G}^1, \dots, \mathcal{G}^T\}$. DySAT applies the attention-based GCN model of Veličković et al. (2018) to each \mathcal{G}^t and obtains $\mathbf{z}_v^1, \dots, \mathbf{z}_v^T$ for every node \mathbf{v} , where \mathbf{z}_v^t encodes the structural information from \mathcal{G}^t . Similar to Vaswani et al. (2017), DySAT adds a positional embedding \mathbf{p}^t to each \mathbf{z}_v^t and obtains $\bar{\mathbf{z}}_v^t = \mathbf{z}_v^t + \mathbf{p}^t$, where each \mathbf{p}^t encodes information about the relative position of the t^{th} snapshot compared to other snapshots. Finally, DySAT applies a multi-head self-attention on $\bar{\mathbf{z}}_v^1, \dots, \bar{\mathbf{z}}_v^T$ to get the final representation of the node to be sent to the decoder.

Kazemi et al. (2019) extend positional encoding to continuous time encoding through a vector representation for time dubbed *Time2Vec*. An interesting direction for future research is to extend DySAT to CTDGs by replacing positional encoding with Time2Vec.

4.6.3. RNN-BASED ENCODERS FOR CTDGS

RNN-based approaches for CTDGs (Kumar et al. (2018b); Trivedi et al. (2017); Ma et al. (2018b); Trivedi et al. (2019)) mainly consider CTDGs where the only possible observation is addition of new edges. They define custom RNNs that update the representations of the source and target nodes forming a new edge (and the representation of the relation between the two nodes in the case of KGs) upon making a new observation ($AddEdge, (v, u), t$) (or ($AddEdge, (v, r, u), t$) in the case of a KG). One of the main differences in these approaches is in the way they define the embedding function and the way they define their custom RNN.

Kumar et al. (2018b), for instance, consider bipartite graphs and develop a model named *JODIE* which defines $EMB(v) = (\mathbf{z}_v, \bar{\mathbf{z}}_v)$ for each node v in the graph where $\mathbf{z}_v \in \mathbb{R}^{d_1}$ and $\bar{\mathbf{z}}_v \in \mathbb{R}^{d_2}$. The values of \mathbf{z}_v are optimized directly (similar to shallow encoders), but the values of $\bar{\mathbf{z}}_v$ are updated using an RNN. In *JODIE*, there are two different RNNs for updating the source and the target nodes. Upon making a new observation ($AddEdge, (v, u), t$), the two RNNs update $\bar{\mathbf{z}}_v$ and $\bar{\mathbf{z}}_u$ as follows:

$$\bar{\mathbf{z}}_v = \text{RNN}_{source}(\bar{\mathbf{z}}_v, [\bar{\mathbf{z}}_u; \Delta t_v; \mathbf{f}]) \quad (41)$$

$$\bar{\mathbf{z}}_u = \text{RNN}_{target}(\bar{\mathbf{z}}_u, [\bar{\mathbf{z}}_v; \Delta t_u; \mathbf{f}]) \quad (42)$$

where Δt_v represents the time elapsed since v 's previous interaction (similarly for Δt_u), \mathbf{f} represents edge features (e.g., it can be the rating a user assigned to a movie), and $[\bar{\mathbf{z}}_v; \Delta t_u; \mathbf{f}]$ represents the concatenation of $\bar{\mathbf{z}}_v$, Δt_u and \mathbf{f} . RNN_{source} is a standard RNN that takes as input the current state $\bar{\mathbf{z}}_v$ and a new input $[\bar{\mathbf{z}}_u; \Delta t_v; \mathbf{f}]$, and outputs an updated state for $\bar{\mathbf{z}}_v$; similarly for RNN_{target} . Note that the two functions are executed simultaneously thus $\bar{\mathbf{z}}_v$ is updated based on the previous value of $\bar{\mathbf{z}}_u$ and $\bar{\mathbf{z}}_u$ is updated based on the previous value of $\bar{\mathbf{z}}_v$. The two vectors \mathbf{z}_v and $\bar{\mathbf{z}}_v$ for a node v are then concatenated as one vector and sent to a decoder for making predictions.

Trivedi et al. (2017) consider KGs and define $EMB(v) = (\mathbf{z}_v)$ for every node v where $\mathbf{z}_v \in \mathbb{R}^{d_1}$ and $EMB(r) = (\mathbf{z}_r)$ for every relation r where $\mathbf{z}_r \in \mathbb{R}^{d_2}$. Their model, named *Know-Evolve*, defines two custom RNNs that update \mathbf{z}_v and \mathbf{z}_u upon making a new observation ($AddEdge, (v, r, u), t$) as follows:

$$\mathbf{z}_v = \text{Tanh}(\mathbf{W}_s \Delta t_v + \mathbf{W}_{hh} \text{Tanh}(\mathbf{W}_h [\mathbf{z}_v; \mathbf{z}_u; \mathbf{r}_{p_v}])) \quad (43)$$

$$\mathbf{z}_u = \text{Tanh}(\mathbf{W}_t \Delta t_u + \mathbf{W}_{hh} \text{Tanh}(\mathbf{W}_h [\mathbf{z}_u; \mathbf{z}_v; \mathbf{r}_{p_u}])) \quad (44)$$

where Δt_v and Δt_u are defined as before, \mathbf{r}_{p_v} is the vector representation for the last relation that v was involved in (similarly for \mathbf{r}_{p_u}), $\mathbf{W}_s \in \mathbb{R}^{d_1 \times 1}$, $\mathbf{W}_t \in \mathbb{R}^{d_1 \times 1}$, $\mathbf{W}_h \in \mathbb{R}^{l \times (2d_1 + d_2)}$, and $\mathbf{W}_{hh} \in \mathbb{R}^{d_1 \times l}$ are weight matrices, and $[\mathbf{z}_u; \mathbf{z}_v; \mathbf{r}_{p_u}]$ is the concatenation of \mathbf{z}_u , \mathbf{z}_v and \mathbf{r}_{p_u} . The vector representation for relations is optimized directly (similar to shallow encoders). Compared to *JODIE*, *Know-Evolve* may be more influenced by Δt s because *Know-Evolve* projects each Δt to a vector and sums the resulting vector with the influence coming from the representations of source, target, and relation. Unlike *JODIE*, the dependence of the update rules on Δt s in *Know-Evolve* is somewhat separate from the two nodes and the relation involved.

Trivedi et al. (2019) develop a model that can be used for several types of graphs. They define $EMB(v) = (\mathbf{z}_v)$. Upon making a new observation ($AddEdge, (v, u), t$), they update the

node representation for v using the following custom RNN (and similarly for u):

$$\mathbf{z}_v = \phi(\mathbf{W}_1 \mathbf{z}_{\mathcal{N}(u)} + \mathbf{W}_2 \mathbf{z}_v + \mathbf{W}_3 \Delta t_v) \quad (45)$$

where $\mathbf{z}_{\mathcal{N}(u)}$ is a weighted aggregation of the neighbors of u , Δt_v is defined as before, \mathbf{W}_i s are weight matrices, and ϕ is an activation function. The aggregation $\mathbf{z}_{\mathcal{N}(u)}$ can be different for different types of graphs (e.g., it can take relations into account in the case of a KG). Trivedi et al. (2019) define a temporal attention mechanism to obtain the neighbor weights for $\mathbf{z}_{\mathcal{N}(u)}$ at each time.

Other ways of defining the embedding function as well as custom RNNs can be viewed in Ma et al. (2018b) for directed graphs, Dai et al. (2016) for bipartite graphs, and Jin et al. (2019) for interpolation in KGs.

4.6.4. DISCUSSION AND ANALYSIS OF RNN-BASED ENCODERS

Information propagation: Upon observing $(AddEdge, (v, u), t)$ (or $(AddEdge, (v, r, u), t)$ in the case of a KG), many existing works only update the nodes directly involved in the new edge. Ma et al. (2018b) argue that it is important to propagate this information to the neighboring nodes so that they can update their representations accordingly. Towards this goal, they first compute a vector representation for the new observation as follows:

$$\mathbf{z}_o = \phi(\mathbf{W}_1 \mathbf{z}_{s_v} + \mathbf{W}_2 \mathbf{z}_{t_u} + \mathbf{b}) \quad (46)$$

where \mathbf{z}_{s_v} and \mathbf{z}_{t_u} belong to $EMB(v)$ and $EMB(u)$ respectively and \mathbf{W}_1 , \mathbf{W}_2 and \mathbf{b} are learnable parameters. \mathbf{z}_o is then sent to the immediate neighbors of v and u and custom RNNs update the representation of the neighbors based on \mathbf{z}_o and based on how they are connected to v and/or u .

Attributed graphs: For attributed graphs where nodes have attributes with fixed values, one way to incorporate these attributes into the model is by initializing (part of) the node representations using their attribute values (Trivedi et al. (2019)). For the case where the attribute values can change over time as well, Seo et al. (2018) and Feng et al. (2018a) develop models that take such changes into account for DTDGs. Developing models for attributed CTDGs where the attribute values can change over time is an interesting direction for future research.

The streaming scenario: For the RNN-based approaches for CTDGs, once the RNN weights are learned during training, the RNN has learned how to take an observation as input and update the node (and relation) embeddings without requiring to compute any further gradients. That is, after training, the RNN weights can be frozen and used for updating the representations as new observations arrive. This makes RNN-based approaches a natural choice for the streaming scenario. Although as the amount of data collected during the test (frozen) time increases (e.g., when it reaches some predefined threshold), the training can run again on all the collected data to learn better weights for the RNN, then the weights can be frozen again and the updated RNN can replace the old one.

4.7. Autoencoder-based Encoders

Consider a DTDG as $\{\mathcal{G}^1, \dots, \mathcal{G}^T\}$ and let $\mathbf{A}^1, \dots, \mathbf{A}^T$ be the corresponding adjacency matrices. Goyal et al. (2017) learn an autoencoder AE^1 for \mathcal{G}^1 similar to SDNE where the

encoder takes as input $\mathbf{A}^1[i]$ and generates a vector representation $\mathbf{z}_{v_i}^1$ for node v_i . The reconstructor takes $\mathbf{z}_{v_i}^1$ as input and reconstructs $\mathbf{A}^1[i]$. $\mathbf{z}_{v_i}^1$ and $\mathbf{z}_{v_j}^1$ are constrained to be close together if there is an edge between v_i and v_j . Having AE^1 , the embedding function for the first snapshot is defined as $\text{EMB}^1(v_i) = (\mathbf{z}_{v_i}^1)$. For the t^{th} snapshot ($t > 1$), an autoencoder AE^t is initialized with the weights from AE^{t-1} and then trained based on \mathbf{A}^t to produce the vector representations for nodes at snapshot t . AE^t can have a different size (e.g., different number of neurons or layers) compared to AE^{t-1} . The authors decide the size of AE^t based on heuristic methods that take into account the size of AE^{t-1} and how different \mathcal{G}^t is from \mathcal{G}^{t-1} . If the size of AE^t is different from AE^{t-1} , in order to still be able to initialize AE^t according to AE^{t-1} , the authors use the Net2WiderNet and Net2DeeperNet approaches from Chen et al. (2015a), which change the number of neurons and the number of layers in an autoencoder without substantially changing the function it computes.

The approach of Goyal et al. (2017) uses the information within previous snapshots of a DTDG to enable learning an autoencoder for the current snapshot faster. Furthermore, initializing AE^t according to AE^{t-1} implicitly acts as a regularizer imposing a smoothness constraint (see Section 4.3). However, the embeddings learned in their approach may not capture the evolution of the nodes. To better capture the evolution of the nodes, Bonner et al. (2018) propose to develop autoencoders that reconstruct a node’s neighborhood in the next snapshot(s) given the current snapshot. They use a two-layer GCN as the encoder and dot-product as the reconstructor of the autoencoder. The authors also propose a variational autoencoder model where instead of directly learning the embeddings \mathbf{Z}^t in the encoder, they learn a Gaussian distribution from which \mathbf{Z}^t is sampled. The Gaussian distribution is parameterized by a mean vector $\boldsymbol{\mu}$ and a variance vector $\boldsymbol{\gamma}$ that are learned using two separate two-layer GCNs with tied parameters on the first layers.

To take more snapshots into account in learning node embeddings, Goyal et al. (2018) propose to learn a single autoencoder where at snapshot t , the encoder takes as input $\mathbf{A}^{t-l}[i], \mathbf{A}^{t-l+1}[i], \dots, \mathbf{A}^{t-1}[i]$ and produces a vector $\mathbf{z}_{v_i}^t$ corresponding to the embedding of v_i at time t , and the reconstructor takes as input $\mathbf{z}_{v_i}^t$ and reconstructs $\mathbf{A}^t[i]$. Goyal et al. (2017) propose several ways for modeling the encoder and reconstructor. Examples for the encoder include feeding a concatenation $[\mathbf{A}^{t-l}[i]; \mathbf{A}^{t-l+1}[i]; \dots; \mathbf{A}^{t-1}[i]]$ into a feed-forward neural network or feeding a sequence $\mathbf{A}^{t-l}[i], \mathbf{A}^{t-l+1}[i], \dots, \mathbf{A}^{t-1}[i]$ into an LSTM. Similar architectures are used for the reconstructor.

Rahman and Al Hasan (2016) also follow an autoencoder-based approach by mapping each node-pair (instead of each node) to a hidden representation, which can then be used to predict addition or deletion of edges in the next snapshot. Towards this goal, for each snapshot they compute features for each pair of nodes based on the graphlet transitions (see Pržulj et al. (2004); Pržulj (2007)). Then they concatenate the features for each node-pair from the previous snapshots (similar to Goyal et al. (2017)) and feed the concatenation to an autoencoder that learns a vector representation for the node-pair.

4.8. Diachronic Encoders

Diachronic encoders are encoders that map every pair (node, timestamp) or every pair (relation, timestamp) to a hidden representation. Note that this is different from Definition 2 where encoders map every node or every relation to a hidden representation. Such encoders

can be used effectively for interpolation as they learn to provide node and relation embeddings at any point in time. Goel et al. (2020) propose a diachronic encoder for nodes of a KG in which each vector $\mathbf{z}_v^t \in \mathbb{R}^d$ in the embedding of v is defined as follows:

$$\mathbf{z}_v^t[i] = \begin{cases} \mathbf{a}_v[i]\phi(\mathbf{w}_v[i]t + \mathbf{b}_v[i]), & \text{if } 1 \leq i \leq \gamma d. \\ \mathbf{a}_v[i], & \text{if } \gamma d < i \leq d. \end{cases} \quad (47)$$

where $\mathbf{a}_v \in \mathbb{R}^d$, $\mathbf{w}_v \in \mathbb{R}^{\gamma d}$ and $\mathbf{b}_v \in \mathbb{R}^{\gamma d}$ are node-specific parameters, γ specifies the percentage of features that are a function of time, and ϕ is an activation function. Notice how the embedding vector \mathbf{z}_v^t is an explicit function of time. To obtain the embedding of a node v at a specific time such as 1999, one can simply replace t in the equation with 1999 and use \mathbf{z}_v^{1999} as the features of v on 1999. While Goel et al. (2020) mainly proposed the above embedding for nodes, it can be also used for relations. Moreover, while they propose this approach for interpolation, it can be potentially adapted to extrapolation.

Xu et al. (2019a) also define \mathbf{z}_v^t to be a direct function of time as follows:

$$\mathbf{z}_v^t = \mathbf{z}_v + \alpha_v \mathbf{w}_v t + \beta_v \sin(2\pi\omega_v t) + \mathcal{N}(0, \Sigma_v) \quad (48)$$

where \mathbf{z}_v is a shallow embedding for v , $\alpha_v \mathbf{w}_v t$ aims at modeling trend, $\beta_v \sin(2\pi\omega_v t)$ aims at modeling seasonality, and $\mathcal{N}(0, \Sigma_v)$ turns the embedding into a Gaussian distribution with mean $\mathbf{z}_v + \alpha_v \mathbf{w}_v t + \beta_v \sin(2\pi\omega_v t)$ and covariance matrix Σ_v .

Dasgupta et al. (2018) develop a diachronic encoder by mapping every timestamp $t \in T$ into a vector representation \mathbf{z}_t (i.e. a shallow encoder for timestamps) and then projecting entity and relation embedding vectors to the space of t as follows:

$$\mathbf{z}_v^t = \mathbf{z}_v - (\mathbf{z}_t' \mathbf{z}_v) \mathbf{z}_t, \quad \mathbf{z}_r^t = \mathbf{z}_r - (\mathbf{z}_t' \mathbf{z}_r) \mathbf{z}_t \quad (49)$$

where \mathbf{z}_v , \mathbf{z}_r and \mathbf{z}_t are node-specific, relation-specific, and timestamp-specific learnable parameters. Note that unlike the encoders in Equation (47) and (48) which provide embeddings at any time t , the above encoder can provide embeddings only for a pre-defined set of timestamps.

Instead of considering shallow embeddings for each timestamp as in Dasgupta et al. (2018), García-Durán et al. (2018) consider shallow embeddings for each character in the timestamps. Then, they map each (relation, timestamp) pair into a vector representation \mathbf{z}_r^t by feeding a shallow embedding of r (i.e. \mathbf{z}_r) and the shallow embeddings of the characters in the timestamp into an LSTM³:

$$\mathbf{z}_r^t = \text{LSTM}(\mathbf{z}_r, \mathbf{z}_{c_1}, \dots, \mathbf{z}_{c_k}) \quad (50)$$

where c_i is the i^{th} character in t . One benefit of this approach is that it can naturally deal with missing values in dates. Note that while García-Durán et al. (2018) mainly proposed the above approach for relation embeddings, it can be used for node embeddings as well.

Diachronic encoders have been also explored for word embeddings to understand how the meaning (or usage) of words has changed over time (see, e.g., Hamilton et al. (2016); Bamler and Mandt (2017)).

3. In cases where there are different time modifiers (e.g., *OccurredAt*, *Since*, and *Until*), they consider a shallow embedding for the time modifier as well and this representation is also fed into the LSTM.

4.9. Staleness

Consider an extrapolation problem over a CTDG and an encoder that updates the embedding for each node v whenever a new observation involving v is made (e.g., when a new edge is added between v and some other node). Assume the last time the encoder updated the embedding for v was at time t_v and currently we are at time t ($> t_v$). Depending on how long it has passed since t_v (corresponding to $t - t_v$), the embedding for v may be staled.

To handle the staleness of representations, Kumar et al. (2018b) propose a method to learn how the representation for a node v evolves when no observation involving v (or involving a node that affects v) is made. Let $\text{EMB}(v) = (\mathbf{z}_v)$ and let $\Delta t_v = t - t_v$. Following the approach proposed by Beutel et al. (2018), Kumar et al. (2018b) first create a vector representation $\mathbf{z}_{\Delta t_v} \in \mathbb{R}^d$ for Δt_v where the i^{th} element of the vector is computed as follows:

$$\mathbf{z}_{\Delta t_v}[i] = \mathbf{w}[i]\Delta t_v + \mathbf{b}[i] \quad (51)$$

where \mathbf{w} and \mathbf{b} are vectors with learnable parameters. Then they compute a new vector representation \mathbf{z}_v^t for v at time t as follows:

$$\mathbf{z}_v^t = (\mathbf{1} + \mathbf{z}_{\Delta t_v}) \odot \mathbf{z}_v \quad (52)$$

where $\mathbf{1} \in \mathbb{R}^d$ is a vector of ones. Having computed \mathbf{z}_v^t , instead of using the (potentially) staled representation \mathbf{z}_v , Kumar et al. (2018b) use \mathbf{z}_v^t to make predictions about v at time t . Note that while diachronic encoders (see Section 4.8) have been mainly proposed for interpolation, if used for extrapolation, several of them (e.g., see Equations (47), (48) and (50)) have the benefit of updating node and relation embeddings even when no observations have been made about the node and relation.

5. Decoders for Dynamic Graphs

We divide the decoders for dynamic graphs into two categories: time-predicting decoders and time-conditioned decoders. In what follows, we explain each category and provide a summary of the existing approaches for that category.

5.1. Time-Predicting Decoders

Time-predicting decoders can be used for extrapolation or interpolation. In extrapolation settings, these decoders aim at predicting *when* an event will happen. For instance, they aim at predicting when *Bob* will visit *Montreal*. In interpolation settings, they aim at predicting a missing timestamp. For instance, they aim at predicting when *Frank Lampard* became the head coach of *Chelsea* assuming $(\text{FrankLampard}, \text{HeadCoachOf}, \text{Chelsea}, ?)$ is in a KG, where $?$ shows that the timestamp is missing.

Sun et al. (2012) were among the first to study the problem of predicting when a particular type of relation will be formed between two nodes. To make such a prediction, first they find all paths between two nodes. These paths are matched with a set of pre-defined path templates and the number of paths matching each template is counted. These counts, which can be roughly considered as node-pair embeddings, are fed into a generalized linear model (GLM) and the score of this model is used to define the parameters of a density function. Sun

et al. (2012) use exponential, Weibull (Weibull et al. (1951)), and geometric distributions for defining the density function. Their exponential distribution for the formation of a relation between two nodes is as follows:

$$\mathbf{f}(t) = \frac{1}{\theta} \exp(-\frac{t}{\theta}) \quad (53)$$

where θ is the output of the GLM model. An expectation of $t \sim \mathbf{f}(t)$ can be used to predict when the relation will be formed between the two nodes, as described in Section 2.7.

Recently there has been growing interest towards time predicting decoders (Trivedi et al. (2019, 2017); Zuo et al. (2018)). Trivedi et al. (2017) consider an encoder that provides an embedding function such that given a dynamic graph until time t gives $\text{EMB}(\mathbf{v}) = (\mathbf{z}_v^t)$ and $\text{EMB}(\mathbf{r}) = (\mathbf{P}_r)$. They first compute a score for the formation of a relation r between two nodes \mathbf{v} and \mathbf{u} as follows:

$$\mathbf{s}_{\mathbf{v},\mathbf{r},\mathbf{u}}(t) = \mathbf{z}_v^{t'} \mathbf{P}_r \mathbf{z}_u^t \quad (54)$$

The obtained score is then used to modulate the conditional intensity function ($\lambda_{\mathbf{v},\mathbf{r},\mathbf{u}}(t|\mathcal{H}_{t-})$) of a TPP for a given relation r and entities \mathbf{v} and \mathbf{u} as follows:

$$\lambda_{\mathbf{v},\mathbf{r},\mathbf{u}}(t|\mathcal{H}_{t-}) = \exp(\mathbf{s}_{\mathbf{v},\mathbf{r},\mathbf{u}}(t))(t - \bar{t}) \quad (55)$$

where \mathcal{H}_{t-} represents the history until time t but not including t , \bar{t} represents the most recent time when either \mathbf{v} or \mathbf{u} was involved in an observation and $t > \bar{t}$. Using \exp ensures that the intensity function is always positive. To predict when relation r will form between \mathbf{v} and \mathbf{u} , one can convert the conditional intensity function into a conditional density function ($\mathbf{f}_{\mathbf{v},\mathbf{r},\mathbf{u}}$) and subsequently take an expectation over the time horizon as described in Section 2.7. Since the intensity function in Equation (55) is piece-wise linear, computing the integral corresponding to the survival function can be done efficiently. During training, however, one needs to compute such an integral $|\mathcal{V}|^2|\mathcal{R}|$ times making the computations intractable. To remedy the intractability, Trivedi et al. (2017) propose an approximation algorithm.

Note that as identified by Jin et al. (2019), the intensity function in Equation (55) may not be suitable for settings involving concurrent events. That is because, in the presence of concurrent events, t may be equal to \bar{t} making the intensity function equivalent to zero.

Trivedi et al. (2019) argue that different types of relations evolve at different rates; e.g., liking posts in a social network occurs more frequently than becoming friends. They model the dynamics of the graph by considering two types of relations: 1- *communications* corresponding to node interactions (e.g., liking someone’s post in social media), 2- *associations* corresponding to topological evolutions (e.g., forming a new friendship). They propose to use different TPPs for these two types of relations. Towards this goal, they assume the embedding function provided by the encoder gives $\text{EMB}(\mathbf{v}) = (\mathbf{z}_v^t)$ and $\text{EMB}(\mathbf{r}) = (\psi_r, \mathbf{z}_r^t)$ and define the intensity function of their TPP as follows:

$$\lambda_{\mathbf{v},\mathbf{r},\mathbf{u}}(t|\mathcal{H}_{t-}) = \psi_r \log(1 + \exp(\frac{\mathbf{z}_r^{t'} [\mathbf{z}_v^t; \mathbf{z}_u^t]}{\psi_r})) \quad (56)$$

where $[\mathbf{z}_v^t; \mathbf{z}_u^t]$ is the concatenation of \mathbf{z}_v^t and \mathbf{z}_u^t . Notice that the above intensity function does not have the $(t - \bar{t})$ term used in Equation (55). Instead, different rates of evolution

(ψ_r) for relations of different types are introduced. Similar to Trivedi et al. (2017), during training Trivedi et al. (2019) also need to compute the integral corresponding to the survival function $|\mathcal{V}|^2|\mathcal{R}|$ times; they approximate the computations through sampling.

Zuo et al. (2018) use the intensity function of a Hawkes process (Hawkes (1971); Mei and Eisner (2017)). The intensity of the interaction is obtained by the Euclidean distance between the interacting nodes and an exponentially discounted interaction history of the neighbors.

For interpolation in KGs, to predict the timestamp for a triple $(v, r, u, ?)$, Leblay and Chekol (2018) and Dasgupta et al. (2018) replace the missing timestamp with all timestamps observed in the KG, find the score for all produced temporal triples using a time-conditioned decoder (see Section 5.2), and select the timestamp resulting in the highest score. Note that this approach may not scale to KGs with many timestamps.

5.2. Time-Conditioned Decoders

Time-conditioned decoders are decoders whose goal is to make predictions for specific timestamps given as input (as opposed to predicting *when* something will happen). These decoders can be used for extrapolation (e.g., predict who will be the CEO of Apple two years from now) or interpolation (e.g., predict who was the CEO of Apple on *2006-04-01*, assuming this piece of information is not explicitly stored in the KG). In other words, time-conditioned decoders predict what happened (or will happen) at some time t where t can be different in different queries. Time-conditioned decoders are mainly similar to static decoders. In what follows, we describe some works that aim at making predictions for a specific timestamp and the decoders they employ.

Dasgupta et al. (2018) develop a model for interpolation in KGs. They use the encoder in Equation (49) to produce node and relation embeddings that are functions of time. To make a prediction about whether some edge (v, r, u) existed at time t or not, they use TransE (see Section 3.2.2 for details) as the decoder:

$$\|\mathbf{z}_v^t + \mathbf{z}_r^t - \mathbf{z}_u^t\| \quad (57)$$

Leblay and Chekol (2018) develop a model for interpolation where a shallow encoder provides node, relation, and timestamp embeddings and the decoder is defined as follows:

$$\|\mathbf{z}_v + \mathbf{z}_r + \mathbf{z}_t - \mathbf{z}_u\| \quad (58)$$

The above decoder can be viewed as an extension of TransE by adding the embedding of the timestamp to the score function. Leblay and Chekol (2018) provide other extensions of TransE as well. Ma et al. (2018c) develop several models for interpolation by using shallow encoders that provide node, relation, and timestamp embeddings and they extend DistMult, ComplEx, RESCAL, and several other decoders by incorporating the timestamp embedding into their score functions. Note that there exists a close connection between these models and the tensor decomposition models (with order 4 tensors) discussed in Section 4.4.

If a shallow encoder is used for the timestamp embeddings (which is the case in the works described so far), then an embedding can only be learned for the timestamps that have been observed in the train set. Therefore, these approaches may not generalize well to the timestamps not observed in the train set as a vector representation has not been learned for

these timestamps. For instance, if a KG does not contain any information with timestamp 08/07/2012, the above approaches do not learn an embedding for this date and may not be able to make predictions about that date. For the same reason, these models cannot be used effectively for predicting something in a future timestamp (i.e. extrapolation) as the training data does not contain any future timestamps. Moreover, these models require many parameters and are prone to overfitting when the number of different timestamps in the training data is large.

The model proposed by García-Durán et al. (2018) addresses the above issues. They use the encoder in Equation (50) to generate an embedding \mathbf{z}_r^t for each relation r at time t ; for nodes, they use a shallow encoder. Then, having the vector representations \mathbf{z}_v , \mathbf{z}_r^t , and \mathbf{z}_u , they use TransE and DistMult as the decoder. Since they use shallow embeddings for each character in the timestamp (not for the timestamp itself), their model can be potentially applied to timestamps unseen during training.

The approaches described so far learn a single static representation (\mathbf{z}_v) for nodes and use this representation to make predictions about the nodes at any time. Goel et al. (2020) argue that learning a static representation for nodes may result in the loss of important information. That is because to make a prediction about a node v in some timestamp t (e.g., predicting the movies v liked on 1990 assuming v is a person), one needs to know the specific properties of v around time t , whereas a static representation only provides an aggregation of v 's properties over time. To address this issue, Goel et al. (2020) use the encoder in Equation (47) to learn node features at any specific time. For the decoder, they use TransE, DistMult, and SimpleE. They prove that using sine as the activation function for Equation (47) and SimpleE as the decoder results in a model that is fully expressive for link prediction for temporal KGs.

Making Predictions for a single timestamp: In cases where all predictions are to be made for a single timestamp or a single time interval, (e.g., predicting what happens in the next snapshot of a DTDG, or predicting what happens in near future *without* predicting when it will happen), the existing approaches mostly use a static decoder from Section 3.2. A notable exception is the work of Zhou et al. (2018b) for link prediction in DTDGs where a point process based on triadic closure is employed. Let v_i , v_j , and v_k be three nodes in a graph at snapshot t . v_i , v_j , and v_k form a *closed triad* if all of them are pair-wise connected, and they form an *open triad* if all but one pair of the nodes are connected. In an open triad, the node that is connected to the other two nodes is called the *center node* of the triad. A fundamental mechanism in the formation and evolution of dynamic networks known as *triad closure* is the process of closed triads being created from open triads (Coleman and Coleman (1994); Huang et al. (2015a)).

Example 12 Consider the graph in Figure 1(b). In this graph, v_1 , v_2 and v_3 form a closed triad. v_2 , v_3 and v_4 form an open triad with v_2 being the center node of this open triad.

For two nodes v and u , let \mathbf{z}_v^t and \mathbf{z}_u^t represent the embedding of the two nodes at the t^{th} snapshot respectively. Zhou et al. (2018b) model the probability of v and u forming an edge in the next snapshot to be proportional to the number of open triads this edge will close and the similarities of \mathbf{z}_v^t and \mathbf{z}_u^t to the embeddings of the center nodes in the open triads involving v and u .

6. Other Relevant Models and Problems

While we mainly provided an overview of the models conforming to an encoder-decoder framework, there are other active lines of work on modeling dynamic (knowledge) graphs. Here, we briefly review some of the other related works and some similar problems.

6.1. Statistical Relational Models

Statistical relational models (Raedt et al. (2016); Koller et al. (2007)) aim at marrying logic and probability to build probabilistic relational models. Examples include Markov logic networks (Richardson and Domingos (2006)), Problog (De Raedt et al. (2007)), probabilistic soft logic (Kimmig et al. (2012)) and relational logistic regression (Kazemi et al. (2014)). These models typically use soft rules such as:

$$\langle w : \text{Friends}(x, y) \wedge \text{Friends}(y, z) \Rightarrow \text{Friends}(x, z) \rangle \quad (59)$$

where the rule implies “friends of friends are likely to be friends” and the weight w of the rule is a measure of confidence for the rule. A model is created using a combination of such soft rules and predictions are made using logical and probabilistic inference. Different models differ in how they interpret these rules and weights. The rules and the weights may be learned from data. One may consider several hidden variables for each node, where each hidden variable corresponds to a feature, and define or learn the correlation between hidden and observed variables using logical rules. The hidden variables in this setting resemble the node representations learned by encoder-decoder architectures with the main difference being that these models learn a set of random variables per node, whereas encoder-decoder architectures learn vectors and matrices of numbers.

In comparison to the encoder-decoder framework, statistical relational models naturally capture uncertainty about facts and relations, which is critical in applications where relations are derived from noisy measurements or ambiguous interpretations such as natural language processing. Furthermore, statistical relational models permit joint inference in a principled and interpretable way over the entire graph while taking into account the uncertainty of the facts. However, this comes at a computational price and therefore it is often needed to restrict the expressivity of the model by only using model structures that are known to be tractable (Van Haaren et al. (2015); Van den Broeck (2011); Kazemi et al. (2016)) or to approximate the inference (Van den Broeck and Darwiche (2013); Kersting et al. (2009); Bui et al. (2013)). In contrast, inference with encoder-decoder models usually scales to large datasets and although the operators used for inference are in some cases questionable, end-to-end learning allows the weights of the operators to be adjusted to yield the best predictions possible for desired tasks.

Statistical relational models have been extended to dynamic cases as well (see, e.g., Gabbay et al. (1998); Sadilek and Kautz (2010); Papai et al. (2012); Dylla et al. (2013); Huber et al. (2014); Chekol et al. (2017); Chekol and Stuckenschmidt (2018)). In their simplest form, these models can be extended to dynamic cases by adding time as an argument similar to the following soft rule:

$$\langle w : \text{Friends}(x, y, t) \wedge \text{Friends}(y, z, t) \wedge \neg \text{Friends}(x, z, t) \Rightarrow \text{Friends}(x, z, t + 1) \rangle \quad (60)$$

where the rule may increase the probability of the worlds where friends of friends become friends in the next snapshot. The amount of increase in the probability is controlled by w . Note how this rule is similar to the triadic closure procedure of Huang et al. (2015a) as it models how closed triads are created from open triads.

Besides the approaches based on soft rules, there exists a family of approaches based on walks where the aim is to learn probabilistic walks on the graphs (or KGs), which, starting from any node v , ends up (probabilistically) at the nodes that have a desired relation with v . These probabilistic walks are different from the random walk approaches discussed in this survey. Examples of such approaches include the works of Lao and Cohen (2010); Lao et al. (2011); Das et al. (2018a). These approaches are quite similar to the approaches based on soft rules corresponding in some cases to a subset of the soft rule approaches (Kazemi and Poole (2018a)).

6.2. Spatiotemporal Graphs

For applications where it is possible to define temporal relations between the nodes, several papers take a DTDG and combine the snapshots through connecting the nodes in different snapshots to get a spatiotemporal graph (ST graph), i.e., a graph that spans both space and time (links across different time steps are known as temporal links whereas links within a time step are known as spatial links). Then, instead of learning on an evolving graph, a model is learned on the resulting (static) ST graph. An ST graph can be considered as a KG with spatial and temporal relations.

Structured time series problems such as video activity recognition and segmentation, and traffic flow prediction are examples of applications that have benefited from a graph-theoretic formulation by creating ST graphs. In the video domain, for instance, a graph is extracted from each video frame and combined with the graphs extracted from other frames. Pandhre et al. (2018) utilize random walk encoders for ST graphs in two ways: 1- creating random walks on the ST graph (which snaps both space and time), 2- first creating spatial random walks at the current snapshot and then temporal random walks over the temporal graph obtained by only keeping the edges between nodes in different snapshots. Their two models show superior results compared to several baselines on temporal prediction problems such as trajectory classification.

For activity recognition, Brendel and Todorovic (2011) learn a structured activity model from the ST graphs obtained from videos. For recognition, they match the ST graph of the given video with the per-class learned activity models. Jain et al. (2016) use RNNs in a structured setting by modeling each node and edge in an ST graph using an RNN. To make the learning feasible, they partition the nodes (and edges) using semantic similarities and share the RNNs amongst the nodes (and edges) within the same partition. Wang and Gupta (2018) pose activity recognition as a graph classification problem. They extract two kinds of graphs from a video, namely the similarity and ST graph. The similarity graph learns connections between objects that are semantically related to each other whereas the ST graph learns connections between objects that overlap in space and time. Then, they utilize graph classification models on the constructed graphs for activity classification. Action recognition has been also modeled as reasoning over a dynamic graph without creating an ST graph. Li et al. (2018) and Ghosh et al. (2018) use skeleton-based datasets for action

recognition. They consider an evolving human skeleton during the course of the action as a DTDG, contrary to the above approaches that use heuristics to combine the individual frame-level graphs. Then they perform graph convolutions in both the temporal as well as the spatial domain. While Li et al. (2018) only use human pose features, Ghosh et al. (2018) also leverage additional contextual cues such as object features, functional relationships, etc.

Another challenging problem formulated as an ST graph is that of traffic flow prediction. Yu et al. (2018a) propose a spatiotemporal convolution block (ST ConvBlock) that consists of temporal gated convolutions and spatial gated convolutions. These ST ConvBlocks are stacked to obtain feature representations for each node in the graph and for the traffic speed prediction. The traffic flow prediction problem has been also modeled as reasoning over the dynamic graph using RNN-based approaches (as discussed in Section 4.6.1) without creating an ST graph (see, e.g., Li et al. (2017b); Yu et al. (2019)).

6.3. Constructing (Dynamic) KGs from Text

For the case of the KGs, while we focused in this survey on methods for using a (dynamic) KG to make predictions about its past, current, or future state, there is a large body of research on how to construct a (dynamic) KG from text (Carlson et al. (2010); Dong et al. (2014); Zhang (2015); Das et al. (2018b)). These approaches typically rely on information extractors to obtain new (probabilistic) facts from text and then add to the KG the new facts that pass a confidence threshold. Besides using the probabilities produced by the information extractors, some works (e.g., Dong et al. (2014)) also leverage the predictive models discussed in this survey to obtain a prior probability for the new facts solely based on what is already in the graph (and not based on textual data). A detailed discussion of these approaches is out of the scope of this paper.

7. Applications, Datasets & Codes

In this section, we provide an overview of the main applications of (dynamic) graphs. We describe some of the datasets that are widely used in the community for representation learning for (knowledge) graphs. We also provide links to online code for some of the works we discussed in the paper.

7.1. Applications

Dynamic Link prediction: A natural problem for a graph is to predict if there is a link (with label r in case of a KG) between two nodes v and u . In the dynamic case, one may be interested in predicting if such a link existed at time t in the past or if it will appear sometime in the future. An example of this would be to predict whether Donald Trump will **visit** China in the next year or not and we query for this link in a KG. Example applications include temporal KG completion, friend recommendation, finding biological connections among species, and predicting obsolete facts in a KG. The most common evaluation metrics for this task include: AOC (area under ROC curve), corresponding to the probability that the predictor gives a higher score to a randomly chosen existing link than a randomly chosen nonexistent one, GMAUC, a geometric mean of AOC, and PRAUC (area under precision-recall curve) - see, for example, Chen et al. (2019a). Another metric named error

rate was considered in Chen et al. (2019a) corresponding to the ratio of the number of mispredicted links to the total number of truly existing links.

Dynamic Entity/relation prediction: One of the fundamental problems in KGs is to predict missing entities or relations. It is a classical problem where we want to predict either a missing head entity $(?, r, u)$, a missing tail entity $(v, r, ?)$, or a missing relation $(v, ?, u)$. In the context of a DTDG, one may want to predict the missing entity or relation in the next snapshot. In the case of a CTDG, one may want to predict the missing entity or relation at a specific timestamp t (e.g., $(?, r, v, t)$). An example mentioned by Trivedi et al. (2017) is to predict who Donald Trump will mention next? Leblay and Chekol (2018) and Dasgupta et al. (2018) considered the task of predicting a missing entity in the temporal case. They rank all entities that can potentially be the missing entity and then find the rank of the actual missing entity. They used mean rank (MR), and the percentage of cases where the actual missing entity is ranked among the top K (known as Hit@ K) to compare the quality of the results. In addition to the above metrics, García-Durán et al. (2018) and Goel et al. (2020) also computed mean reciprocal rank (MRR). MRR is generally reported under two settings: raw and filtered (see Bordes et al. (2013) for the details).

Recommender systems: The design of dynamic recommender systems is an important applied dynamic graph problem faced by a myriad of e-commerce companies in online retail, video streaming, and music streaming, to name a few examples. In a dynamic recommender-system problem, we have a set of users, a set of items, and a set of timestamped interactions between users and items, and we seek to recommend items to users based on their current tastes (Wu et al. (2017); Kumar et al. (2018b)). The key is that tastes may exhibit cyclic or trend behavior which should be adapted to or anticipated inasmuch as this is possible. As a coarse approximation, we can view a dynamic recommender-system problem as one of dynamic link prediction (or dynamic entity prediction) and attempt to recommend to users the items they are likely to autonomously choose. However, a more fine-grained analysis reveals several complications present in recommender systems that may be absent in other dynamic link prediction problems. First, the actual output of a recommender system for a given user is a sequence of slates of recommended items. If the items on an output slate are highly similar, then their utility to the user may be strongly correlated, so that there is a non-negligible risk of the slate being useless to the user. This risk is mitigated with a more diverse slate, even if the diverse slate has a lower sum of expected utilities than the uniform slate (see Kaminskis and Bridge (2017))⁴. Second, from the point of view of an e-commerce company, the purpose of a recommender system is typically to maximize profit in the long term. Therefore, it may be desirable to recommend to the user items in which the user has no immediate interest, but which are expected to cause the user to purchase profitable items or click on profitable advertisements. From this perspective, the design of a dynamic recommender system can be seen as a reinforcement learning problem on a dynamic graph. One seeks to dynamically control an evolving graph to maximize some profit-related objective.

Time Prediction: For dynamic graphs, besides predicting which event will happen in the future, an interesting problem is to also predict when that event will happen. As compared to other tasks, this task only exists for dynamic networks. For instance, we saw the example

4. This reasoning resembles Markowitz’s portfolio theory (Markowitz (1952)).

in Section 5.1 of predicting when *Bob* will visit *Montreal*. A similar time prediction problem is the temporal scoping problem in KG completion where the goal is to predict missing timestamps (e.g., answering queries such as $(v, r, u, ?)$, where v is known to have had a relation r with u in the past). Sun et al. (2012) and Trivedi et al. (2017, 2019) used the mean absolute error between the predicted time and ground truth to measure the quality of the results. Dasgupta et al. (2018) ordered the predicted timestamps in the decreasing order of their probabilities and selected the rank associated with the correct timestamp. They computed the mean rank (MR) to compare the results.

Node classification: Node classification is the problem of classifying graph nodes into different classes. An example of a node classification problem is to predict the political affiliation of the users of a social network based on their attributes, connections, and activities. In particular, one may be interested in making such predictions in a streaming scenario where the classification scores keep updating as new events happen or as a user activity is observed. Node classification is often studied under two settings: transductive and inductive. In the transductive setting (also known as semi-supervised classification), given the labels of a few nodes, we want to predict the labels of the other nodes in the graph. In the inductive setting, the label is to be predicted for new nodes that have not been seen during training. The problem of node classification becomes challenging in the dynamic case as the distribution of the class labels may change over time. There are not many publicly-available datasets for dynamic node classification. Pareja et al. (2019) used Elliptic a network of bitcoin transactions, for temporal node classification. Sato et al. (2019) considered a dataset of face to face proximity for temporal node classification. They used the SIR model for node classification, which is a popular framework to model the spread of epidemics (Hethcote (2000)). At each timestamp, each node can be one of three possible states: susceptible (S), infectious (I), and recovered (R). Classification accuracy is a widely used metric for this task. For datasets where the task is a multi-class classification, micro-F1 and macro-F1 scores are also used to measure the performance (Cui et al. (2018)).

Graph classification: Graph classification is the problem of classifying the whole graph into one class from a set of predefined classes. This task can be useful in domains like bioinformatics and social networks. In bioinformatics, for instance, one important application is the protein function classification where proteins are viewed as graphs. As another example, Taheri et al. (2019) modeled activity state classification for a troop of GPS-tracked baboons as a dynamic graph classification problem where the labels are different activity states such as sleeping, hanging-out, coordinated non-progression, and coordinated progression. Common graph classification benchmarks include COLLAB (Yanardag and Vishwanathan (2015)), PROTEINS (Borgwardt et al. (2005)) and D&D (Dobson and Doig (2003)). The performance is typically measured in terms of classification accuracy.

Network clustering: Network clustering or detecting communities in graphs such as social networks, biological networks, etc. is an important problem. A network cluster/community typically refers to a subset of nodes that are densely connected, but loosely connected to the rest of the nodes. One challenge in community detection for dynamic graphs is that one needs to model how communities evolve. A common performance measure used for this task is the overlap between the predicted and the true cluster assignments. Xin et al. (2016) utilize the random walk encoders to find the closely associated nodes for a given node and cluster the global network into overlapping communities. Furthermore, the closely associated nodes

are updated when impacted by dynamic events, giving a dynamic community detection method. Chen et al. (2019b) recast community detection as a node-wise classification problem and present a family of graph neural networks for solving the community detection problem in a supervised setting. Crawford and Milenković (2018) considered a dynamic communication network between 75 patients and health workers in France with four different clusters representing doctors, administrators, nurses, and patients.

Dynamic Question/query answering: The way search engines respond to our questions has evolved in the last few years. While traditionally search engines were aiming at suggesting documents in which the answer to a query question can be potentially found, these days they try to directly answer the question. This has become possible in part due to question answering over KGs (QA-KG). Jia et al. (2018) further extend this work to temporal questions. They defined a temporal question as a question that has a temporal expression (date, time, interval, and periodic events) or temporal signal (before, after, during, etc.) in the question or whose answer is temporal. For QA-KG, the most popular metrics are precision, recall, F1 score, AUC, and accuracy. An example of a temporal question given by Jia et al. (2018) was, “Which teams did Neymar play for before joining PSG?” Hamilton et al. (2018) propose a way of mapping a query formulated as a conjunctive first-order logic formula into a vector representation (corresponding to a query embedding) through geometric operations. The geometric operations leading to query embeddings are jointly optimized with node embeddings such that the query embedding is close in the embedded space to the embeddings of the entities corresponding to the correct answer of the query. For QA-KG, the most popular metrics are precision, recall, F1 score, AUC, and accuracy.

7.2. Datasets

There is a large collection of datasets used for research on static and dynamic graphs in the community. For brevity, we will only survey a representative sample of these datasets. To explore more network datasets, we refer readers to several popular network repositories such as Stanford Large Network Dataset Collection (<https://snap.stanford.edu/data/index.html>), Network Repository (<http://networkrepository.com/index.php>), Social Computing data repository (<http://socialcomputing.asu.edu/pages/datasets>), LINQS (<https://linqs.soe.ucsc.edu/data>), UCI Network Data Repository (<https://networkdata.ics.uci.edu/>), CNetS Data Repository (<http://cnets.indiana.edu/resources/data-repository/>) and Koblenz Network Collection (<http://konect.uni-koblenz.de/networks/>). Section 7 of Cui et al. (2018) and Section 7.1 of Zhang et al. (2018b) are also good starting points to explore other datasets.

Table 2 gives a brief summary of the datasets. In Section 7.2.1, we describe some evolving and temporal datasets. In Section 7.2.2, we survey some of the popular traditional datasets that have been widely used in the community but are not necessarily temporal or dynamic.

-
5. <http://snap.stanford.edu/data/#as>
 6. <https://linqs.soe.ucsc.edu/node/236>
 7. <http://networkrepository.com/radoslaw-email.php>
 8. <https://dblp.uni-trier.de/xml>
 9. <https://snap.stanford.edu/data/com-DBLP.html>
 10. <https://www.kaggle.com/ellipticco/elliptic-data-set>

Dataset	Type	Nodes	Edges	Granularity
Social Evolution	Social network	83	376-791 Associations 2,016,339 Communications	6 mins
Github	Social network	12,328	70,640-166,565 Associations 604,649 Communications	-
HEP-TH	Citation network	1,424-7,980	2,556-21,036	Monthly
Autonomous systems	Communication network	103-6,474 ⁵	243-13,233	Daily
GDELT	Events knowledge graph	14,018	31.29M	15 mins
ICEWS	Events knowledge graph	12498	0.67M	Daily
YAGO	Knowledge graph	15,403	138,056	Mostly yearly
Wikidata	Knowledge graph	11,134	150,079	Yearly
Reddit	Social network	55,863	858,49	Seconds
Enron	Email network	151	50.5K	Seconds
FB-Forum	Social network	899	33.7K	Seconds
Blog	Social network	5,196	171,743	-
Cora ⁶	Citation network	2708	5429	-
Flicker	Social network	1,715,256	22,613,981	-
UCI	Communication network	1,899	59,835	Seconds
Radoslaw ⁷	Email network	167	82.9K	Seconds
DBLP ⁸⁹	Citation network	315,159	743,70	-
YELP	Bipartite ratings	6,569	95,361	Seconds
MovieLens-10M	Bipartite ratings	20,537	43,760	Seconds
CONTACT	Face-to-face proximity	274	28,200	Seconds
HYPERTEXT09	Face-to-face proximity	113	20,800	Seconds
Elliptic ¹⁰	Bitcoin transactions	203,769	234,355	49 time steps

Table 2: A summary of the datasets used in dynamic (knowledge) graph publications, the type of data they contain, the number of nodes and edges they contain, and their temporal granularity.

7.2.1. TEMPORAL DATASETS

Social Evolution Dataset¹¹: The social evolution dataset was released by the MIT Human Dynamics Lab (Madan et al. (2012)) and is used by Trivedi et al. (2019). The dataset is collected between Jan 2008 to Sep 2008 and has 83 nodes. As mentioned in Section 5.1, Trivedi et al. (2019) consider two categories of relations: associations (or topological evolution) and communications (or interactions). The number of associations evolves from 376 initial edges to 791 final edges. The number of communication events (proximity, calls, and SMS) in the dataset is 2,016,339.

Github Dataset: Github is a web-based hosting service for codes. Trivedi et al. (2019) collected a dataset from Github archives between Jan 2013 and Dec 2013. They consider “following a user” as an associative event and “starring” or “watching” a repository as a communicative event. The dataset has 12,328 nodes. The number of associations evolves from 70,640 initial edges to 166,565 final edges. There are 604,649 communication events between the users in this dataset.

11. <http://realitycommons.media.mit.edu/socialrevolution.html>

HEP-TH: Gehrke et al. (2003) created a dataset of arXiv papers in the High Energy Physics Theory conference from January 1993 to April 2003. The graph is a citation network where nodes represent papers and directed edges represent the citations. Goyal et al. (2017) made this an evolving graph by considering all published papers up to that month. Their graph evolves from 1424 nodes to 7980 nodes, and from 2,556 edges to 21,036 edges. Goyal et al. (2018) and Yu et al. (2018b) also conducted experiments on variants of this dataset.

Autonomous Systems¹²: Autonomous systems graph (Leskovec et al. (2005)) is a communication network from the Border Gateway Protocol logs. The graph has 733 daily snapshots from Nov 1997 to Jan 2000. The graph grows from 103 to 6,474 nodes and from 243 to 13,233 edges. One unique thing about this dataset is that while most other graphs have only the addition of the nodes and edges, this graph has instances of both addition and deletion of nodes and edges. Goyal et al. (2017) use the first 100 snapshots, and Goyal et al. (2018) use the last 50 snapshots for their experiments.

GDELT: Global Database of Events, Language, and Tone (GDELT) (Leetaru and Schrodt (2013)) is an initiative to construct a database of all the events across the globe connecting people, organizations, events, news sources, and locations. Trivedi et al. (2017) collected a subset of this data from April 1, 2015, to Mar 31, 2016, with the temporal granularity of 15 mins. This subset contains 14,018 nodes, 20 types of relations, and 31.29M edges. Goel et al. (2020) use this dataset with modified train, validation, and test sets for temporal KG completion.

ICEWS¹³: Integrated Crisis Early Warning System (ICEWS) dataset (Boschee et al. (2015)) contains information about political events with their timestamps. Here the entities represent important political people (presidents, prime ministers, bureaucrats) and countries, and the relations are political scenarios such as negotiate, sign a formal agreement, criticize, etc. The data¹⁴ used by (Trivedi et al. (2017)) is collected from Jan 1, 2014 to Dec 31, 2014 with the temporal granularity of 24 hours and has 12,498 nodes, 260 types of relations, and 0.67M facts. García-Durán et al. (2018) created two KGs based on the ICEWS dataset. One of these KGs¹⁵ contains information from 2014 and has 6,869 nodes, 230 types of relations, and 96,730 facts. The other one is a longer-term KG¹⁶ containing the events occurring between 2005-2015 with 10,094 nodes, 251 types of relations, and 461,329 facts.

YAGO: YAGO¹⁷ (Hoffart et al. (2013)) is a spatially and temporally enriched version of the Wikipedia knowledge base. The nodes represent people, groups, artifacts and events while the relations represent facts such as *wasBornIn*, *playsFor*, *isLocatedIn*, etc. YAGO contains temporal information in the form of “occursSince” and “occursUntil”. The dataset¹⁸ created by García-Durán et al. (2018) has 15,403 nodes, 34 types of relations and 138,056 facts. Jiang et al. (2016)’s dataset has 9,513 nodes, 10 types of relations, and 15,914 facts.

12. <http://snap.stanford.edu/data/as-733.html>

13. <http://www.icews.com/>

14. <https://github.com/rstriv/Know-Evolve/tree/master/data/icews>

15. <https://github.com/nle-ml/mmkb/tree/master/TemporalKGs/icews14>

16. <https://github.com/nle-ml/mmkb/tree/master/TemporalKGs/icews05-15>

17. <https://www.mpi-inf.mpg.de/departments/databases-and-information-systems/research/yago-naga/yago/>

18. <https://github.com/nle-ml/mmkb/tree/master/TemporalKGs/yago15k>

Wikidata: Leblay and Chekol (2018) created a temporal knowledge base¹⁹ using Wikidata. García-Durán et al. (2018) considered a subset of this dataset by selecting the most frequent entities along with the relations that include these entities. In their dataset²⁰, they have 11,134 nodes, 95 types of relations, and 150,079 facts.

Reddit HyperLink Network²¹: Subreddit hyperlink network (Kumar et al. (2018a)) is a directed network extracted from the posts that create hyperlinks from one subreddit to another. Each edge has temporal information, the sentiment of the source towards the target, and the text of the source post. The dataset also comes with subreddit embeddings for 51,278 subreddits. There are 55,863 nodes and 858,490 edges in the graph.

Enron²²: Enron email dataset (Klimt and Yang (2004)) is the network of email exchanges among the employees of Enron. This data was originally released by the Federal Energy Regulatory Commission as part of their investigation. There are several variants²³ of this dataset available and it has been widely used in the community (Nguyen et al. (2018b); De Winter et al. (2018); Chen et al. (2018a); Sankar et al. (2018)).

FB-FORUM²⁴: This dataset comes from a Facebook-like online community of students at the University of California, Irvine and was collected in 2004 (Opsahl (2011); Rossi and Ahmed (2015)). This is a bipartite graph where the nodes represent students and groups while the edges represent students’ broadcast messages on the groups. The dataset used by Nguyen et al. (2018b) for their experiments had 899 nodes and 33.7K edges along with timestamp in Unix time.

UCI²⁵: This dataset (Kunegis (2013)) is obtained from the same social network as the one for FB-FORUM. The dataset is a communication network among users along with timestamps. This network has 1,899 nodes and 59,835 edges (Ma et al. (2018b); Sankar et al. (2018); Yu et al. (2018b)).

YELP²⁶: The YELP dataset is a subset of YELP’s businesses, reviews, and user data. It was originally made public as a Kaggle contest. The dataset consists of 6,685,900 reviews by 1,637,138 users for 192,609 businesses. Sankar et al. (2018) use a part of the YELP dataset where they select the businesses in the state of Arizona and retain businesses that have more than 15 reviews.

MovieLens-10M²⁷: The MovieLens dataset (Harper and Konstan (2016)) is a dynamic user-tag interactions dataset. It consists of 10 million ratings and 100,000 tag applications applied to 10,000 movies by 72,000 users. The dataset shows the tagging behavior of users on the movies they rated. Sankar et al. (2018) utilize a subset of this dataset with 20,537 nodes and 43,760 links.

CONTACT: The CONTACT dataset introduced by (Chaintreau et al. (2007)) is a dynamic network for face-to-face proximity collected through wireless devices carried by

19. <https://staff.aist.go.jp/julien.leblay/datasets/>

20. <https://github.com/nle-ml/mmkb/tree/master/TemporalKGs/wikidata>

21. <http://snap.stanford.edu/data/soc-RedditHyperlinks.html>

22. https://en.wikipedia.org/wiki/Enron_Corpus

23. <https://snap.stanford.edu/data/email-Enron.html>, <https://www.cs.cmu.edu/~enron/>, <http://networkrepository.com/>, <https://www.kaggle.com/wcukierski/enron-email-dataset>

24. <http://networkrepository.com/fb-forum.php>

25. <http://konect.uni-koblenz.de/networks/opsahl-ucsocial>

26. <https://www.yelp.com/dataset>

27. <https://grouplens.org/datasets/movielens/10m/>

Reference	Code	Nature	Type
Jin et al. (2019)	https://github.com/INK-USC/RE-Net	Dynamic	KG
Goel et al. (2020)	https://github.com/borealisai/DE-Simple	Dynamic	KG
Wu et al. (2019)	https://github.com/lienwc/DKGE/	Dynamic	KG
Trivedi et al. (2017)	https://github.com/rstriv/Know-Evolve	Dynamic	KG
Dasgupta et al. (2018)	https://github.com/malllabiisc/HyTE	Dynamic	KG
Bian et al. (2019)	https://github.com/Change2vec/change2vec	Dynamic	KG
Seo et al. (2018)	https://github.com/youngjoo-epfl/gconvRNN	Dynamic	G
Wu et al. (2017)	https://github.com/RuidongZ/Recurrent_Recommender_Networks	Dynamic	G
Zhou et al. (2018b)	https://github.com/luckiezhou/DynamicTriad	Dynamic	G
Zhang et al. (2018d)	https://github.com/ZW-ZHANG/TIMERS	Dynamic	G
Sajjad et al. (2019)	https://github.com/shps/incremental-representation-learning	Dynamic	G
Kazemi and Poole (2018c)	https://github.com/Mehran-k/Simple	Static	KG
Dong et al. (2017)	https://ericdongyx.github.io/metapath2vec/m2v.html	Static	KG
Zhang et al. (2018a)	https://github.com/daokunzhang/MetaGraph2Vec	Static	KG
Trouillon et al. (2016)	https://github.com/ttrouill/ComplEx	Static	KG
Lacroix et al. (2018)	https://github.com/facebookresearch/kbc	Static	KG
Nickel et al. (2016b)	https://github.com/mnick/holographic-embeddings	Static	KG
Lerer et al. (2019)	https://github.com/facebookresearch/PyTorch-BigGraph	Static	KG
Kipf and Welling (2017)	https://github.com/tkipf/gcn	Static	G
Hamilton et al. (2017a)	https://github.com/williamleif/GraphSAGE	Static	G
Chen et al. (2015a)	https://github.com/soumith/net2net_torch	Static	G
Grover and Leskovec (2016)	https://github.com/aditya-grover/node2vec	Static	G
Veličković et al. (2018)	https://github.com/PetarV-/GAT	Static	G
Liao et al. (2019)	https://github.com/lrjconan/LanczosNetwork	Static	G

Table 3: Link to open-source software for static and dynamic (knowledge) graphs. “G” indicates that the code can handle only simple homogeneous graphs, whereas “KG” indicates that the code can handle knowledge graphs.

people; a link is created between two people whenever they interact. It contains data from 274 people and 28.2K interactions as described in Chen et al. (2018a).

HYPERTEXT09: The HYPERTEXT09 dataset described in Isella et al. (2011) is a proximity network of attendees at the ACM Hypertext 2009 conference. The dataset contains 113 nodes each corresponding to an attendee and 20,800 edges each corresponding to an interaction between two attendees.

7.2.2. STATIC DATASETS

Blog²⁸: This dataset was collected from the Blog Catalog website. Bloggers follow other bloggers forming graph edges and they categorize their blog under some predefined classes. The graph used by Liu et al. (2019) has 5,196 nodes, 171,743 edges, and 6 types of node classes. Other works using this dataset include Li et al. (2017a) and Ma et al. (2018a).

CiteSeer²⁹: This is a citation network where papers are considered as nodes and citations are considered as the edges. The broad category of papers is used as the class labels. It has 3,312 nodes, 4,732 edges, and 6 types of node classes. This dataset is widely used including in Liu et al. (2019), Kipf and Welling (2017) and Liao et al. (2019).

Flickr^{30 31}: The Flickr dataset (Tang and Liu (2009)) is obtained from a network of users on a photo-sharing website. There are class labels that correspond to the groups that users have subscribed to on the website. The instance used by Tang et al. (2015) contains 1,715,256 nodes and 22,613,981 edges.

28. <http://socialcomputing.asu.edu/datasets/BlogCatalog>

29. <https://linqs.soe.ucsc.edu/node/236>

30. <http://socialnetworks.mpi-sws.org/data-imc2007.html>

31. <https://snap.stanford.edu/data/web-flickr.html>

7.3. Open-Source Software

There are several open-source libraries providing implementations for the papers discussed in the survey. Fey and Lenssen (2019) provide an implementation for many GCN-based papers (e.g., Kipf and Welling (2017); Hamilton et al. (2017a)) in PyTorch (Paszke et al. (2017)). Wang and Ga (2018) also provide implementations for a variety of graph-based models. Along with graph-structured models, they also make tensor-based models like transformer (Vaswani et al. (2017)) available with the intention of facilitating the development of new models combining the two categories. They provide both PyTorch and MXNet (Chen et al. (2015b)) backends for this library. Apart from these libraries, the implementation of several techniques covered in the survey is available in independent code repositories. Table 3 provides links to where these implementations can be found.

8. Future Directions & Conclusion

A wide range of real-world problems can be formulated as reasoning over graphs (or networks). Traditionally, graph analytic methods have been mostly focused on static graphs, while in a large number of applications the graphs are dynamic and evolve. In the past few years, there has been a surge of works on dynamic graphs. In this paper, we surveyed the recent approaches for representation learning over dynamic graphs. We described these approaches according to an encoder-decoder framework, a framework that has gained popularity within several communities.

Our survey sheds light on several ways in which learning from and reasoning with dynamic graphs can be done. Here, we mention some directions to improve learning and reasoning with dynamic graphs.

- Current representation learning algorithms have been mostly designed for discrete-time dynamic graphs (DTDGs), with only a few works on learning from continuous-time dynamic graphs (CTDGs). Even the few existing works for CTDGs are quite limited in the types of observations they can handle as they mainly handle the addition of new edges. A promising direction for future research is to extend the existing models for representation learning over CTDGs, or develop new ones, to deal with other types of observations such as edge deletion, node addition, node deletion, node splitting, node merging, etc. Ma et al. (2018a) take some initial steps towards handling node addition, but their proposal provides an embedding for a new node considering only the current state of the graph (not the evolution of the graph).
- While some of the existing encoders work with certain types of graphs, it is not trivial how they can be used for other types of graphs. For instance, random walk approaches have been mainly designed for non-attributed graphs. Extending these approaches to the case of attributed graphs is not straightforward. The same goes for autoencoder-based encoders where it is not trivial how these approaches can be extended to KGs, and several other models discussed in the survey. An interesting direction for future research would be to extend these models to be applicable to more types of graphs.

- The approaches for CTDGs that can be used for the streaming scenario are mainly based on RNNs. In other sequence modeling domains (e.g., natural language processing), new sequence modeling approaches have been developed some times showing superior performance compared to RNNs. One example is the transformer architecture (Vaswani et al. (2017)) where recurrence has been replaced with self-attention. Designing new models for the streaming scenario in CTDGs based on self-attentions is a promising direction for future research. Other possibilities include designing new models based on neural ordinary differential equations (see, e.g., Chen et al. (2018b)).
- In a CTDG, many observations may be made at the same time. For instance, in an email communications network, a sender may send an email to many receivers at the same time. Current RNN-based encoders consider a random ordering for such observations. This naive approach may hinder an RNN from learning to generalize to other possible orderings of these simultaneous observations. A future direction would be to extend RNN approaches for dealing with simultaneous observations.
- While there exists a body of research on classifying static graphs and some of them may be (to some extent) applicable to dynamic graphs, the literature on classifying dynamic graphs is still at its infancy. When the classification of a dynamic graph is required (e.g., for activity recognition from videos), current approaches often convert the dynamic graph into a static graph and then run a graph classification algorithm on the static graph. Designing dynamic graph classification models is an interesting direction for future research.
- Expressiveness is an important property to be taken into account when selecting/designing a model. A model that is not expressive enough is doomed to underfitting at least for some applications. While a few recent works study the expressiveness of some models for (knowledge) graphs (Trouillon et al. (2017); Kazemi and Poole (2018c); Xu et al. (2019b); Morris et al. (2019)), a more detailed and in-depth study of the expressiveness and its empirical impact may be a promising direction for future research.
- The existing models for dynamic graphs only consider edges connecting two nodes in the graph. However, in real-world applications, some edges may connect more than two nodes. These edges are known as hyperedges. In a KG, for instance, an edge corresponding to a *purchase* relation may connect a person as the *buyer* to another person as the *seller* and also to an item that is being purchased. Kazemi (2018) argues that representation learning algorithms may fail if these hyperedges are converted into (several) binary edges through reifying new entities as, during test time, an embedding does not exist for the reified entities. Some recent works study ways of handling such hyperedges for static (knowledge) graphs (Wen et al. (2016); Feng et al. (2018b); Yadati et al. (2018); Bai et al. (2019); Fatemi et al. (2019b)). A future research direction would be to develop models for dynamic (knowledge) graphs that are capable of handling hyperedges.
- Recently published papers on modeling dynamic graphs are each tested on different datasets, making it difficult to compare these models. It would be quite helpful to

create some standard benchmarks with train, validation, and test splits so future models can be tested on the same benchmarks and splits.

References

- Charu Aggarwal and Karthik Subbian. Evolutionary network analysis: A survey. *ACM Computing Surveys (CSUR)*, 47(1):10, 2014.
- Nahla Mohamed Ahmed and Ling Chen. An efficient algorithm for link prediction in temporal uncertain social networks. *Information Sciences*, 331:120–136, 2016.
- Nahla Mohamed Ahmed, Ling Chen, Yulong Wang, Bin Li, Yun Li, and Wei Liu. Sampling-based algorithm for link prediction in temporal networks. *Information Sciences*, 374:1–14, 2016.
- Hirotsugu Akaike. Fitting autoregressive models for prediction. *Annals of the institute of Statistical Mathematics*, 21(1):243–247, 1969.
- Mohammad Al Hasan and Mohammed J Zaki. A survey of link prediction in social networks. In *Social network Data Analytics*, pages 243–275. Springer, 2011.
- László Babai and Ludik Kucera. Canonical labelling of graphs in linear average time. In *20th Annual Symposium on Foundations of Computer Science (sfcs 1979)*, pages 39–46. IEEE, 1979.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- Song Bai, Feihu Zhang, and Philip HS Torr. Hypergraph convolution and hypergraph attention. *arXiv preprint arXiv:1901.08150*, 2019.
- Ivana Balazević, Carl Allen, and Timothy M Hospedales. Hypernetwork knowledge graph embeddings. *arXiv preprint arXiv:1808.07018*, 2018.
- Ivana Balažević, Carl Allen, and Timothy M Hospedales. Tucker: Tensor factorization for knowledge graph completion. In *ICML Workshop on Adaptive & Multitask Learning: Algorithms & Systems*, 2019.
- Robert Bamler and Stephan Mandt. Dynamic word embeddings. In *ICML*, pages 380–389, 2017.
- Nikolaos Bastas, Theodoros Semertzidis, Apostolos Axenopoulos, and Petros Daras. evolve2vec: Learning network representations using temporal unfolding. In *International Conference on Multimedia Modeling*, pages 447–458. Springer, 2019.
- Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps and spectral techniques for embedding and clustering. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 585–591, 2001.

- Alex Beutel, Paul Covington, Sagar Jain, Can Xu, Jia Li, Vince Gatto, and Ed H Chi. Latent cross: Making use of context in recurrent recommender systems. In *Eleventh ACM International Conference on Web Search and Data Mining*, pages 46–54. ACM, 2018.
- Ranran Bian, Yun Sing Koh, Gillian Dobbie, and Anna Divoli. Network embedding and change modeling in dynamic heterogeneous networks. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 861–864. ACM, 2019.
- Cemal Cagatay Bilgin and Bülent Yener. Dynamic network evolution: Models, clustering, anomaly detection. *IEEE Networks*, 2006.
- Stephen Bonner, John Brennan, Ibad Kureshi, Georgios Theodoropoulos, Andrew Stephen McGough, and Boguslaw Obara. Temporal graph offset reconstruction: Towards temporally robust graph representation learning. In *2018 IEEE International Conference on Big Data (Big Data)*, pages 3737–3746. IEEE, 2018.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. Translating embeddings for modeling multi-relational data. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2787–2795, 2013.
- Karsten M Borgwardt, Cheng Soon Ong, Stefan Schönauer, SVN Vishwanathan, Alex J Smola, and Hans-Peter Kriegel. Protein function prediction via graph kernels. *Bioinformatics*, 21(suppl_1):i47–i56, 2005.
- Elizabeth Boschee, Jennifer Lautenschlager, Sean O’Brien, Steve Shellman, James Starz, and Michael Ward. Icews coded event data. *Harvard Dataverse*, 12, 2015.
- George EP Box, Gwilym M Jenkins, Gregory C Reinsel, and Greta M Ljung. *Time series analysis: forecasting and control*. John Wiley & Sons, 2015.
- Matthew Brand. Fast low-rank modifications of the thin singular value decomposition. *Linear Algebra and its Applications*, 415(1):20–30, 2006.
- William Brendel and Sinisa Todorovic. Learning spatiotemporal graphs of human activities. In *Thirteenth International Conference on Computer Vision (ICCV)*, pages 778–785. IEEE, 2011.
- Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.
- Hung Hai Bui, Tuyen N Huynh, and Sebastian Riedel. Automorphism groups of graphical models and lifted variational inference. In *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence (UAI)*, 2013.
- J.R. Bunch and C.P. Nielsen. Updating the singular value decomposition. *Numerische Mathematik*, 31:111 – 129, 1978.

- Hongyun Cai, Vincent W Zheng, and Kevin Chen-Chuan Chang. A comprehensive survey of graph embedding: Problems, techniques, and applications. *IEEE Transactions on Knowledge and Data Engineering*, 30(9):1616–1637, 2018.
- Jin-Yi Cai, Martin Fürer, and Neil Immerman. An optimal lower bound on the number of variables for graph identification. *Combinatorica*, 12(4):389–410, 1992.
- Shaosheng Cao, Wei Lu, and Qiongkai Xu. Deep neural networks for learning graph representations. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- Andrew Carlson, Justin Betteridge, Bryan Kisiel, Burr Settles, Estevam R Hruschka, and Tom M Mitchell. Toward an architecture for never-ending language learning. In *Twenty-Fourth AAAI Conference on Artificial Intelligence*, 2010.
- Augustin Chaintreau, Pan Hui, Jon Crowcroft, Christophe Diot, Richard Gass, and James Scott. Impact of human mobility on opportunistic forwarding algorithms. *IEEE Transactions on Mobile Computing*, pages 606–620, 2007.
- Deepayan Chakrabarti, Ravi Kumar, and Andrew Tomkins. Evolutionary clustering. In *Twelfth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 554–560. ACM, 2006.
- Benjamin Paul Chamberlain, James Clough, and Marc Peter Deisenroth. Neural embeddings of graphs in hyperbolic space. *arXiv preprint arXiv:1705.10359*, 2017.
- Chris Chatfield and Mohammad Yar. Holt-winters forecasting: Some practical issues. *Journal of the Royal Statistical Society. Series D (The Statistician)*, 37(2), 1988.
- Melisachew Wudage Chekol and Heiner Stuckenschmidt. Rule based temporal inference. In *Technical Communications of the 33rd International Conference on Logic Programming (ICLP 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- Melisachew Wudage Chekol, Giuseppe Pirrò, Joerg Schoenfish, and Heiner Stuckenschmidt. Marrying uncertainty and time in knowledge graphs. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- Chen Chen and Hanghang Tong. Fast eigen-functions tracking on dynamic graphs. *Society for Industrial and Applied Mathematics Publications*, pages 559–567, 2015.
- Jinyin Chen, Xuanheng Xu, Yangyang Wu, and Haibin Zheng. GC-LSTM: Graph convolution embedded LSTM for dynamic link prediction. *arXiv preprint arXiv:1812.04206*, 2018a.
- Jinyin Chen, Jian Zhang, Xuanheng Xu, Chengbo Fu, Dan Zhang, Qingpeng Zhang, and Qi Xuan. E-LSTM-D: A deep learning framework for dynamic network link prediction. *arXiv preprint arXiv:1902.08329*, 2019a.
- Tian Qi Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 6572–6583, 2018b.

- Tianqi Chen, Ian Goodfellow, and Jonathon Shlens. Net2net: Accelerating learning via knowledge transfer. *arXiv preprint arXiv:1511.05641*, 2015a.
- Tianqi Chen, Mu Li, Yutian Li, Min Lin, Naiyan Wang, Minjie Wang, Tianjun Xiao, Bing Xu, Chiyuan Zhang, and Zheng Zhang. Mxnet: A flexible and efficient machine learning library for heterogeneous distributed systems. *arXiv preprint arXiv:1512.01274*, 2015b.
- Ting Chen and Yizhou Sun. Task-guided and path-augmented heterogeneous network embedding for author identification. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, pages 295–304. ACM, 2017.
- Zhengdao Chen, Lisha Li, and Joan Bruna. Supervised community detection with line graph neural networks. In *International Conference on Learning Representations (ICLR)*, 2019b.
- Zhengdao Chen, Soledad Villar, Lei Chen, and Joan Bruna. On the equivalence between graph isomorphism testing and function approximation with GNNs. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019c.
- Yun Chi, Xiaodan Song, Dengyong Zhou, Koji Hino, and Belle L Tseng. On evolutionary spectral clustering. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 3(4):17, 2009.
- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- Edward Choi, Mohammad Taha Bahadori, Andy Schuetz, Walter F Stewart, and Jimeng Sun. Doctor AI: Predicting clinical events via recurrent neural networks. In *Machine Learning for Healthcare Conference*, pages 301–318, 2016.
- James S Coleman and James Samuel Coleman. *Foundations of social theory*. Harvard University Press, 1994.
- David R Cox and Peter Adrian Walter Lewis. Multivariate point processes. In *Sixth Berkeley Symposium on Mathematical Statistics and Probability*, volume 3, pages 401–448, 1972.
- Joseph Crawford and Tijana Milenković. Cluenet: Clustering a temporal network based on topological similarity rather than denseness. *PloS one*, 13(5):e0195993, 2018.
- Peng Cui, Xiao Wang, Jian Pei, and Wenwu Zhu. A survey on network embedding. *IEEE Transactions on Knowledge and Data Engineering*, 2018.
- Paulo Ricardo da Silva Soares and Ricardo Bastos Cavalcante Prudêncio. Time series based link prediction. In *International Joint Conference on Neural Networks (IJCNN)*, pages 1–7. IEEE, 2012.
- Hanjun Dai, Yichen Wang, Rakshit Trivedi, and Le Song. Deep coevolutionary network: Embedding user and item features for recommendation. *arXiv preprint arXiv:1609.03675*, 2016.

- Hanjun Dai, Zornitsa Kozareva, Bo Dai, Alex Smola, and Le Song. Learning steady-states of iterative algorithms over graphs. In *International Conference on Machine Learning (ICML)*, pages 1114–1122, 2018.
- Rajarshi Das, Shehzaad Dhuliawala, Manzil Zaheer, Luke Vilnis, Ishan Durugkar, Akshay Krishnamurthy, Alex Smola, and Andrew McCallum. Go for a walk and arrive at the answer: Reasoning over paths in knowledge bases using reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2018a.
- Rajarshi Das, Tsendsuren Munkhdalai, Xingdi Yuan, Adam Trischler, and Andrew McCallum. Building dynamic knowledge graphs from text using machine reading comprehension. *arXiv preprint arXiv:1810.05682*, 2018b.
- Shib Sankar Dasgupta, Swayambhu Nath Ray, and Partha Talukdar. Hyte: Hyperplane-based temporally aware knowledge graph embedding. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 2001–2011, 2018.
- Chandler Davis and W. M. Kahan. The rotation of eigenvectors by a perturbation. iii. *SIAM Journal of Numerical Analysis*, 7:1–46, 1970.
- Luc De Raedt, Angelika Kimmig, and Hannu Toivonen. Problog: A probabilistic prolog and its application in link discovery. In *Twentieth International Joint Conference on Artificial Intelligence (IJCAI)*, volume 7, pages 2462–2467. Hyderabad, 2007.
- Sam De Winter, Tim Decuyper, Sandra Mitrović, Bart Baesens, and Jochen De Weerd. Combining temporal aspects of dynamic networks with node2vec for a more efficient dynamic link prediction. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 1234–1241. IEEE, 2018.
- Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 3844–3852, 2016.
- Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Łukasz Kaiser. Universal transformers. *arXiv preprint arXiv:1807.03819*, 2018.
- Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. Convolutional 2d knowledge graph embeddings. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Paul D Dobson and Andrew J Doig. Distinguishing enzyme structures from non-enzymes without alignments. *Journal of Molecular Biology*, 330(4):771–783, 2003.
- Xin Dong, Evgeniy Gabrilovich, Jeremy Heitz, Wilko Horn, Ni Lao, Kevin Murphy, Thomas Strohmman, Shaohua Sun, and Wei Zhang. Knowledge vault: A web-scale approach to probabilistic knowledge fusion. In *Twentieth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 601–610. ACM, 2014.

- Yuxiao Dong, Nitesh V Chawla, and Ananthram Swami. metapath2vec: Scalable representation learning for heterogeneous networks. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 135–144. ACM, 2017.
- Lun Du, Yun Wang, Guojie Song, Zhicong Lu, and Junshan Wang. Dynamic network embedding: An extended approach for skip-gram based network embedding. In *Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2086–2092, 2018.
- Nan Du, Hanjun Dai, Rakshit Trivedi, Utkarsh Upadhyay, Manuel Gomez-Rodriguez, and Le Song. Recurrent marked temporal point processes: Embedding event history to vector. In *Twenty-Second ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1555–1564. ACM, 2016.
- Daniel M Dunlavy, Tamara G Kolda, and Evrim Acar. Temporal link prediction using matrix and tensor factorizations. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 5(2):10, 2011.
- David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P Adams. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2224–2232, 2015.
- Maximilian Dylla, Iris Miliaraki, and Martin Theobald. A temporal-probabilistic database model for information extraction. *Very Large Data Bases (VLDB)*, 6(14):1810–1821, 2013.
- Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(211), 1936.
- Jeffrey L Elman. Finding structure in time. *Cognitive Science*, 14(2):179–211, 1990.
- Robert F Engle and Jeffrey R Russell. Autoregressive conditional duration: a new model for irregularly spaced transaction data. *Econometrica*, pages 1127–1162, 1998.
- Cristóbal Esteban, Volker Tresp, Yinchong Yang, Stephan Baier, and Denis Krompaß. Predicting the co-evolution of event and knowledge graphs. In *2016 19th International Conference on Information Fusion (FUSION)*, pages 98–105. IEEE, 2016.
- Bahare Fatemi, Siamak Ravanbakhsh, and David Poole. Improved knowledge graph embedding using background taxonomic information. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3526–3533, 2019a.
- Bahare Fatemi, Perouz Taslakian, David Vazquez, and David Poole. Knowledge hypergraphs: Prediction beyond binary relations. *arXiv preprint arXiv:1906.00137*, 2019b.
- Fuli Feng, Xiangnan He, Xiang Wang, Cheng Luo, Yiqun Liu, and Tat-Seng Chua. Temporal relational ranking for stock prediction. *arXiv preprint arXiv:1809.09441*, 2018a.

- Yifan Feng, Haoxuan You, Zizhao Zhang, Rongrong Ji, and Yue Gao. Hypergraph neural networks. *arXiv preprint arXiv:1809.09401*, 2018b.
- Matthias Fey and Jan E. Lenssen. Fast graph representation learning with PyTorch Geometric. *CoRR*, abs/1903.02428, 2019.
- Matthias Fey, Jan Eric Lenssen, Frank Weichert, and Heinrich Müller. Splinecnn: Fast geometric deep learning with continuous b-spline kernels. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 869–877, 2018.
- Dov M Gabbay, Christopher John Hogger, and John Alan Robinson. *Handbook of Logic in Artificial Intelligence and Logic Programming: Volume 5: Logic Programming*. Clarendon Press, 1998.
- Alberto García-Durán, Sebastijan Dumančić, and Mathias Niepert. Learning sequence encoders for temporal knowledge graph completion. *arXiv preprint arXiv:1809.03202*, 2018.
- Johannes Gehrke, Paul Ginsparg, and Jon Kleinberg. Overview of the 2003 kdd cup. *ACM SIGKDD Explorations Newsletter*, 5(2):149–151, 2003.
- Pallabi Ghosh, Yi Yao, Larry S Davis, and Ajay Divakaran. Stacked spatio-temporal graph convolutional networks for action segmentation. *arXiv preprint arXiv:1811.10575*, 2018.
- Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International Conference on Machine Learning (ICML)*, pages 1263–1272. JMLR. org, 2017.
- Rishab Goel, Seyed Mehran Kazemi, Marcus Brubaker, and Pascal Poupert. Diachronic embedding for temporal knowledge graph completion. In *Thirty-Fourth AAAI Conference on Artificial Intelligence*, 2020.
- Palash Goyal, Nitin Kamra, Xinran He, and Yan Liu. Dyngem: Deep embedding method for dynamic graphs. In *IJCAI International Workshop on Representation Learning for Graphs*, 2017.
- Palash Goyal, Sujit Rokka Chhetri, and Arquimedes Canedo. dyngraph2vec: Capturing network dynamics using dynamic graph representation learning. *arXiv preprint arXiv:1809.02657*, 2018.
- Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. In *Twenty-Second ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 855–864. ACM, 2016.
- Ekta Gujral, Ravdeep Pasricha, and Evangelos E Papalexakis. SamBaTen: Sampling-based batch incremental tensor decomposition. In *SIAM International Conference on Data Mining*, pages 387–395. SIAM, 2018.
- İsmail Güneş, Şule Gündüz-Öğüdücü, and Zehra Çataltepe. Link prediction using time series of neighborhood-based node similarity scores. *Data Mining and Knowledge Discovery*, 30(1):147–180, 2016.

- Manish Gupta, Charu C Aggarwal, Jiawei Han, and Yizhou Sun. Evolutionary clustering and analysis of bibliographic networks. In *2011 International Conference on Advances in Social Networks Analysis and Mining*, pages 63–70. IEEE, 2011.
- Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1024–1034, 2017a.
- Will Hamilton, Payal Bajaj, Marinka Zitnik, Dan Jurafsky, and Jure Leskovec. Embedding logical queries on knowledge graphs. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2026–2037, 2018.
- William L Hamilton, Jure Leskovec, and Dan Jurafsky. Diachronic word embeddings reveal statistical laws of semantic change. In *Association for Computational Linguistics (ACL)*, 2016.
- William L Hamilton, Rex Ying, and Jure Leskovec. Representation learning on graphs: Methods and applications. *IEEE Data Engineering Bulletin*, 40(3):52–74, 2017b.
- F Maxwell Harper and Joseph A Konstan. The movielens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems (TIIS)*, 5(4):19, 2016.
- Richard A. Harshman, Peter Ladefoged, H. Graf von Reichenbach, Robert I. Jennrich, Dale Terbeek, Lee Cooper, Andrew L. Comrey, Peter M. Bentler, Jeanne Yamane, and Diane Vaughan. Foundations of the parafac procedure: Models and conditions for an "explanatory" multimodal factor analysis. *UCLA Working Papers Phonetics*, 6:1–84, 1970.
- Alan G Hawkes. Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58(1):83–90, 1971.
- Katsuhiko Hayashi and Masashi Shimbo. On the equivalence of holographic and complex embeddings for link prediction. *arXiv preprint arXiv:1702.05563*, 2017.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. Teaching machines to read and comprehend. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1693–1701, 2015.
- Herbert W Hethcote. The mathematics of infectious diseases. *SIAM review*, 42(4):599–653, 2000.
- Geoffrey E Hinton and Ruslan R Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006.
- Ryohei Hisano. Semi-supervised graph embedding approach to dynamic link prediction. In *International Workshop on Complex Networks*, pages 109–121. Springer, 2018.
- Frank L Hitchcock. The expression of a tensor or a polyadic as a sum of products. *Journal of Mathematics and Physics*, 6(1-4):164–189, 1927.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.

- Johannes Hoffart, Fabian M Suchanek, Klaus Berberich, and Gerhard Weikum. Yago2: A spatially and temporally enhanced knowledge base from wikipedia. *Artificial Intelligence*, 194:28–61, 2013.
- Hao Hu and Guo-Jun Qi. State-frequency memory recurrent neural networks. In *International Conference on Machine Learning (ICML)*, pages 1568–1577, 2017.
- Hong Huang, Jie Tang, Lu Liu, JarDer Luo, and Xiaoming Fu. Triadic closure pattern analysis and prediction in social networks. *IEEE Transactions on Knowledge and Data Engineering*, 27(12):3374–3389, 2015a.
- Zan Huang and Dennis KJ Lin. The time-series link prediction problem with applications in communication surveillance. *INFORMS Journal on Computing*, 21(2):286–303, 2009.
- Zhiheng Huang, Wei Xu, and Kai Yu. Bidirectional LSTM-CRF models for sequence tagging. *arXiv preprint arXiv:1508.01991*, 2015b.
- Jakob Huber, Christian Meilicke, and Heiner Stuckenschmidt. Applying markov logic for debugging probabilistic temporal knowledge bases. In *Fourth Workshop on Automated Knowledge Base Construction (AKBC)*, 2014.
- Nahla Mohamed Ahmed Ibrahim and Ling Chen. Link prediction in dynamic social networks by integrating different types of information. *Applied Intelligence*, 42(4):738–750, 2015.
- Lorenzo Isella, Juliette Stehlé, Alain Barrat, Ciro Cattuto, Jean-François Pinton, and Wouter Van den Broeck. What’s in a crowd? analysis of face-to-face behavioral networks. *Journal of Theoretical Biology*, 271(1):166–180, 2011.
- Valerie Isham and Mark Westcott. A self-correcting point process. *Stochastic Processes and Their Applications*, 8(3):335–347, 1979.
- Ashesh Jain, Amir R Zamir, Silvio Savarese, and Ashutosh Saxena. Structural-RNN: Deep learning on spatio-temporal graphs. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 5308–5317, 2016.
- Zhen Jia, Abdalghani Abujabal, Rishiraj Saha Roy, Jannik Strötgen, and Gerhard Weikum. Tequila: Temporal question answering over knowledge bases. In *Twenty-Seventh ACM International Conference on Information and Knowledge Management*, pages 1807–1810. ACM, 2018.
- Tingsong Jiang, Tianyu Liu, Tao Ge, Lei Sha, Baobao Chang, Sujian Li, and Zhifang Sui. Towards time-aware knowledge graph completion. In *Twenty-Sixth International Conference on Computational Linguistics: Technical Papers*, pages 1715–1724, 2016.
- Woojeong Jin, Changlin Zhang, Pedro Szekely, and Xiang Ren. Recurrent event network for reasoning over temporal knowledge graphs. *arXiv preprint arXiv:1904.05530*, 2019.
- Marius Kaminskis and Derek Bridge. Diversity, serendipity, novelty, and coverage: A survey and empirical analysis of beyond-accuracy objectives in recommender systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 7(1):2, 2017.

- Seyed Mehran Kazemi. *Representing and learning relations and properties under uncertainty*. PhD thesis, University of British Columbia, 2018.
- Seyed Mehran Kazemi and David Poole. Bridging weighted rules and graph random walks for statistical relational models. *Frontiers in Robotics and AI*, 5:8, 2018a.
- Seyed Mehran Kazemi and David Poole. RelNN: A deep neural model for relational learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018b.
- Seyed Mehran Kazemi and David Poole. Simple embedding for link prediction in knowledge graphs. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 4289–4300, 2018c.
- Seyed Mehran Kazemi, David Buchman, Kristian Kersting, Sriraam Natarajan, and David Poole. Relational logistic regression. In *Fourteenth International Conference on the Principles of Knowledge Representation and Reasoning*, 2014.
- Seyed Mehran Kazemi, Angelika Kimmig, Guy Van den Broeck, and David Poole. New liftable classes for first-order probabilistic inference. In *Advances in Neural Information Processing Systems*, pages 3117–3125, 2016.
- Seyed Mehran Kazemi, Rishab Goel, Sepehr Eghbali, Janahan Ramanan, Jaspreet Sahota, Sanjay Thakur, Stella Wu, Cathal Smyth, Pascal Poupart, and Marcus Brubaker. Time2vec: Learning a vector representation of time. *arXiv preprint arXiv:1907.05321*, 2019.
- Nicolas Keriven and Gabriel Peyré. Universal invariant and equivariant graph neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- Kristian Kersting, Babak Ahmadi, and Sriraam Natarajan. Counting belief propagation. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 277–284. AUAI Press, 2009.
- Min-Soo Kim and Jiawei Han. A particle-and-density based evolutionary clustering method for dynamic networks. *Very Large Data Bases (VLDB)*, 2(1):622–633, 2009.
- Angelika Kimmig, Stephen H Bach, Matthias Broecheler, Bert Huang, and Lise Getoor. A short introduction to probabilistic soft logic. In *NeurIPS Workshop on Probabilistic Programming: Foundations and Applications*, volume 1, page 3, 2012.
- John Frank Charles Kingman. Poisson processes. *Encyclopedia of biostatistics*, 6, 2005.
- Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017.
- Bryan Klimt and Yiming Yang. Introducing the enron corpus. In *CEAS*, 2004.
- Daphne Koller, Nir Friedman, Sašo Džeroski, Charles Sutton, Andrew McCallum, Avi Pfeffer, Pieter Abbeel, Ming-Fai Wong, David Heckerman, Chris Meek, et al. *Introduction to statistical relational learning*. MIT press, 2007.

- Jacob Krantz and Jugal Kalita. Abstractive summarization using attentive neural techniques. *arXiv preprint arXiv:1810.08838*, 2018.
- Srijan Kumar, William L Hamilton, Jure Leskovec, and Dan Jurafsky. Community interaction and conflict on the web. In *World Wide Web Conference on World Wide Web*, pages 933–943. International World Wide Web Conferences Steering Committee, 2018a.
- Srijan Kumar, Xikun Zhang, and Jure Leskovec. Learning dynamic embedding from temporal interaction networks. *arXiv preprint arXiv:1812.02289*, 2018b.
- Jérôme Kunegis. Konect: The koblenz network collection. In *Twenty-Second International Conference on World Wide Web, WWW '13 Companion*, pages 1343–1350, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-2038-2. doi: 10.1145/2487788.2488173. URL <http://doi.acm.org/10.1145/2487788.2488173>.
- Timothée Lacroix, Nicolas Usunier, and Guillaume Obozinski. Canonical tensor decomposition for knowledge base completion. In *International Conference on Machine Learning (ICML)*, 2018.
- Ni Lao and William W Cohen. Relational retrieval using a combination of path-constrained random walks. *Machine Learning*, 81(1):53–67, 2010.
- Ni Lao, Tom Mitchell, and William W Cohen. Random walk inference and learning in a large scale knowledge base. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 529–539. Association for Computational Linguistics, 2011.
- Julien Leblay and Melisachew Wudage Chekol. Deriving validity time in knowledge graph. In *Companion of the The Web Conference 2018 on The Web Conference 2018*, pages 1771–1776. International World Wide Web Conferences Steering Committee, 2018.
- Kalev Leetaru and Philip A Schrodtt. Gdelt: Global data on events, location, and tone, 1979–2012. In *ISA Annual Convention*, volume 2, pages 1–49. Citeseer, 2013.
- Adam Lerer, Ledell Wu, Jiajun Shen, Timothee Lacroix, Luca Wehrstedt, Abhijit Bose, and Alex Peysakhovich. Pytorch-biggraph: A large-scale graph embedding system. *arXiv preprint arXiv:1903.12287*, 2019.
- Jure Leskovec, Jon Kleinberg, and Christos Faloutsos. Graphs over time: densification laws, shrinking diameters and possible explanations. In *Eleventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 177–187. ACM, 2005.
- Pierre-David Letourneau, Muthu Baskaran, Tom Henretty, James Ezick, and Richard Lethin. Computationally efficient cp tensor decomposition update framework for emerging component discovery in streaming data. In *2018 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–8, 09 2018. doi: 10.1109/HPEC.2018.8547700.
- Ron Levie, Federico Monti, Xavier Bresson, and Michael M Bronstein. Caylennets: Graph convolutional neural networks with complex rational spectral filters. *IEEE Transactions on Signal Processing*, 67(1):97–109, 2017.

- Keith Levin, Farbod Roosta-Khorasani, Michael W Mahoney, and Carey E Priebe. Out-of-sample extension of graph adjacency spectral embedding. *arXiv preprint arXiv:1802.06307*, 2018.
- Chaolong Li, Zhen Cui, Wenming Zheng, Chunyan Xu, and Jian Yang. Spatio-temporal graph convolution for skeleton based action recognition. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Jundong Li, Harsh Dani, Xia Hu, Jiliang Tang, Yi Chang, and Huan Liu. Attributed network embedding for learning in a dynamic environment. In *Conference on Information and Knowledge Management (CIKM)*, pages 387–396. ACM, 2017a.
- Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. *arXiv preprint arXiv:1707.01926*, 2017b.
- Yang Li, Nan Du, and Samy Bengio. Time-dependent representation for neural event sequence prediction. *arXiv preprint arXiv:1708.00065*, 2017c.
- Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493*, 2015.
- Renjie Liao, Zhizhen Zhao, Raquel Urtasun, and Richard S Zemel. Lanczosnet: Multi-scale deep graph convolutional networks. *arXiv preprint arXiv:1901.01484*, 2019.
- David Liben-Nowell and Jon Kleinberg. The link-prediction problem for social networks. *Journal of the American Society for Information Science and Technology*, 58(7):1019–1031, 2007.
- Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. Learning entity and relation embeddings for knowledge graph completion. In *Twenty-ninth AAAI conference on artificial intelligence*, 2015.
- Hanxiao Liu, Yuexin Wu, and Yiming Yang. Analogical inference for multi-relational embeddings. In *International Conference on Machine Learning (ICML)*, pages 2168–2178. JMLR. org, 2017.
- Xi Liu, Ping-Chun Hsieh, Nick Duffield, Rui Chen, Muhe Xie, and Xidao Wen. Real-time streaming graph embedding through local actions. In *Companion Proceedings of the 2019 World Wide Web Conference (WWW '19 Companion)*, 2019.
- Jianxin Ma, Peng Cui, and Wenwu Zhu. Depthlgp: learning embeddings of out-of-sample nodes in dynamic networks. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018a.
- Yao Ma, Ziyi Guo, Zhaochun Ren, Eric Zhao, Jiliang Tang, and Dawei Yin. Streaming graph neural networks. *arXiv preprint arXiv:1810.10627*, 2018b.
- Yunpu Ma, Volker Tresp, and Erik A Daxberger. Embedding models for episodic knowledge graphs. *Journal of Web Semantics*, 2018c.

- Anmol Madan, Manuel Cebrian, Sai Moturu, Katayoun Farrahi, et al. Sensing the” health state” of a community. *IEEE Pervasive Computing*, 11(4):36–45, 2012.
- Sedigheh Mahdavi, Shima Khoshraftar, and Aijun An. dynnode2vec: Scalable dynamic network embedding. In *2018 IEEE International Conference on Big Data (Big Data)*, pages 3762–3765. IEEE, 2018.
- Franco Manessi, Alessandro Rozza, and Mario Manzo. Dynamic graph convolutional networks. *arXiv preprint arXiv:1704.06199*, 2017.
- Harry Markowitz. Portfolio selection. *The journal of finance*, 7(1):77–91, 1952.
- Haggai Maron, Ethan Fetaya, Nimrod Segol, and Yaron Lipman. On the universality of invariant networks. In *ICML*, 2019.
- Hongyuan Mei and Jason M Eisner. The neural hawkes process: A neurally self-modulating multivariate point process. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 6754–6764, 2017.
- Grégoire Mesnil, Yann Dauphin, Kaisheng Yao, Yoshua Bengio, Li Deng, Dilek Hakkani-Tur, Xiaodong He, Larry Heck, Gokhan Tur, Dong Yu, et al. Using recurrent neural networks for slot filling in spoken language understanding. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23(3):530–539, 2015.
- Tomáš Mikolov, Martin Karafiát, Lukáš Burget, Jan Černocký, and Sanjeev Khudanpur. Recurrent neural network based language model. In *Eleventh Annual Conference of the International Speech Communication Association*, 2010.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 3111–3119, 2013.
- Amin Milani Fard, Ebrahim Bagheri, and Ke Wang. Relationship prediction in dynamic heterogeneous information networks. In *European Conference on Information Retrieval*, pages 19–34. Springer, 2019.
- Sudatta Mohanty and Alexey Pozdnukhov. Graph CNN+LSTM framework for dynamic macroscopic traffic congestion prediction. In *International Workshop on Mining and Learning with Graphs*, 2018.
- Behnaz Moradabadi and Mohammad Reza Meybodi. A novel time series link prediction method: Learning automata approach. *Physica A*, 2017.
- Christopher Morris, Martin Ritzert, Matthias Fey, William L Hamilton, Jan Eric Lenssen, Gaurav Rattan, and Martin Grohe. Weisfeiler and leman go neural: Higher-order graph neural networks. In *Thirty-Third AAAI Conference on Artificial Intelligence*, 2019.
- Kevin Patrick Murphy and Stuart Russell. Dynamic bayesian networks: representation, inference and learning. 2002.

- Apurva Narayan and Peter HO'N Roe. Learning graph dynamics using deep neural networks. *IFAC-PapersOnLine*, 51(2):433–438, 2018.
- Deepak Nathani, Jatin Chauhan, Charu Sharma, and Manohar Kaul. Learning attention-based embeddings for relation prediction in knowledge graphs. *arXiv preprint arXiv:1906.01195*, 2019.
- Daniel Neil, Michael Pfeiffer, and Shih-Chii Liu. Phased LSTM: Accelerating recurrent network training for long or event-based sequences. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 3882–3890, 2016.
- Dat Quoc Nguyen, Kairit Sirts, Lizhen Qu, and Mark Johnson. Stranse: a novel embedding model of entities and relationships in knowledge bases. In *NAACL-HLT*, 2016.
- Giang H Nguyen, John Boaz Lee, Ryan A Rossi, Nesreen K Ahmed, Eunyee Koh, and Sungchul Kim. Dynamic network embeddings: From random walks to temporal random walks. In *2018 IEEE International Conference on Big Data (Big Data)*, pages 1085–1092. IEEE, 2018a.
- Giang Hoang Nguyen, John Boaz Lee, Ryan A Rossi, Nesreen K Ahmed, Eunyee Koh, and Sungchul Kim. Continuous-time dynamic network embeddings. In *Companion of the The Web Conference 2018 on The Web Conference 2018*, pages 969–976. International World Wide Web Conferences Steering Committee, 2018b.
- Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. A three-way model for collective learning on multi-relational data. In *International Conference on Machine Learning (ICML)*, volume 11, pages 809–816, 2011.
- Maximilian Nickel, Kevin Murphy, Volker Tresp, and Evgeniy Gabrilovich. A review of relational machine learning for knowledge graphs. *IEEE*, 104(1):11–33, 2016a.
- Maximilian Nickel, Lorenzo Rosasco, and Tomaso Poggio. Holographic embeddings of knowledge graphs. In *Thirtieth AAAI conference on artificial intelligence*, 2016b.
- Mathias Niepert, Mohamed Ahmed, and Konstantin Kutzkov. Learning convolutional neural networks for graphs. In *International Conference on Machine Learning (ICML)*, pages 2014–2023, 2016.
- T. Opsahl. Triadic closure in two-mode networks: Redefining the global and local clustering coefficients. *Social Networks*, 2011.
- Mingdong Ou, Peng Cui, Jian Pei, Ziwei Zhang, and Wenwu Zhu. Asymmetric transitivity preserving graph embedding. In *Twenty-Second ACM SIGKDD Conderence on Knowledge Discovery and Data Mining*, 2016.
- Supriya Pandhre, Himangi Mittal, Manish Gupta, and Vineeth N Balasubramanian. Stwalk: learning trajectory representations in temporal graphs. In *ACM India Joint International Conference on Data Science and Management of Data*, pages 210–219. ACM, 2018.

- Tivadar Papai, Henry Kautz, and Daniel Stefankovic. Slice normalized dynamic markov logic networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1907–1915, 2012.
- Aldo Pareja, Giacomo Domeniconi, Jie Chen, Tengfei Ma, Toyotaro Suzumura, Hiroki Kanezashi, Tim Kaler, and Charles E Leiserson. Evolvegn: Evolving graph convolutional networks for dynamic graphs. *arXiv preprint arXiv:1902.10191*, 2019.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NeurIPS 2017 Autodiff Workshop: The Future of Gradient-based Machine Learning Software and Techniques*, 2017.
- Yulong Pei, Jianpeng Zhang, GH Fletcher, and Mykola Pechenizkiy. Node classification in dynamic social networks. *Proceedings of AALTD*, page 54, 2016.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, 2014.
- Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. Deepwalk: Online learning of social representations. In *Twentieth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 701–710. ACM, 2014.
- Bryan Perozzi, Vivek Kulkarni, and Steven Skiena. Walklets: Multiscale graph embeddings for interpretable network classification. *arXiv preprint arXiv:1605.02115*, 2016.
- Nataša Pržulj. Biological network comparison using graphlet degree distribution. *Bioinformatics*, 23(2):e177–e183, 2007.
- Nataša Pržulj, Derek G Corneil, and Igor Jurisica. Modeling interactome: scale-free or geometric? *Bioinformatics*, 20(18):3508–3515, 2004.
- Stephan Rabanser, Oleksandr Shchur, and Stephan Günnemann. Introduction to tensor decompositions and their applications in machine learning. *arXiv preprint arXiv:1711.10781*, 2017.
- Lawrence R Rabiner and Biing-Hwang Juang. An introduction to hidden markov models. *iee assp magazine*, 3(1):4–16, 1986.
- Luc De Raedt, Kristian Kersting, Sriraam Natarajan, and David Poole. Statistical relational artificial intelligence: Logic, probability, and computation. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 10(2):1–189, 2016.
- Mahmudur Rahman and Mohammad Al Hasan. Link prediction in dynamic networks using graphlet. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 394–409. Springer, 2016.
- Matthew Richardson and Pedro Domingos. Markov logic networks. *Machine Learning*, 62(1-2):107–136, 2006.

- Petar Ristoski and Heiko Paulheim. Rdf2vec: Rdf graph embeddings for data mining. In *International Semantic Web Conference*, pages 498–514. Springer, 2016.
- Ryan A. Rossi and Nesreen K. Ahmed. The network data repository with interactive graph analytics and visualization. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015. URL <http://networkrepository.com>.
- Adam Sadilek and Henry Kautz. Recognizing multi-agent activities from gps data. In *Twenty-Fourth AAAI Conference on Artificial Intelligence*, 2010.
- Hooman Peiro Sajjad, Andrew Docherty, and Yuriy Tyshetskiy. Efficient representation learning using random walks for dynamic graphs. *arXiv preprint arXiv:1901.01346*, 2019.
- Aravind Sankar, Yanhong Wu, Liang Gou, Wei Zhang, and Hao Yang. Dynamic graph representation learning via self-attention networks. *arXiv preprint arXiv:1812.09430*, 2018.
- Adam Santoro, David Raposo, David G Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Timothy Lillicrap. A simple neural network module for relational reasoning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 4967–4976, 2017.
- Purnamrita Sarkar, Sajid M Siddiqi, and Geogrey J Gordon. A latent space approach to dynamic embedding of co-occurrence data. In *Artificial Intelligence and Statistics*, pages 420–427, 2007.
- Koya Sato, Mizuki Oka, Alain Barrat, and Ciro Cattuto. Dyane: Dynamics-aware node embedding for temporal networks, 2019.
- Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. Modeling relational data with graph convolutional networks. In *European Semantic Web Conference*, pages 593–607. Springer, 2018.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. Bidirectional attention flow for machine comprehension. *arXiv preprint arXiv:1611.01603*, 2016.
- Youngjoo Seo, Michaël Defferrard, Pierre Vandergheynst, and Xavier Bresson. Structured sequence modeling with graph convolutional recurrent networks. In *International Conference on Neural Information Processing*, pages 362–373. Springer, 2018.
- Umang Sharan and Jennifer Neville. Temporal-relational classifiers for prediction in evolving domains. In *2008 Eighth IEEE International Conference on Data Mining*, pages 540–549. IEEE, 2008.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. Self-attention with relative position representations. In *Annual Conference of the North American Chapter of the Association for Computational Linguistics*, 2018.
- Chuan Shi, Yitong Li, Jiawei Zhang, Yizhou Sun, and S Yu Philip. A survey of heterogeneous information network analysis. *IEEE Transactions on Knowledge and Data Engineering*, 29(1):17–37, 2016.

- Chuan Shi, Binbin Hu, Wayne Xin Zhao, and S Yu Philip. Heterogeneous information network embedding for recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 31(2):357–370, 2018.
- Martin Simonovsky and Nikos Komodakis. Dynamic edge-conditioned filters in convolutional neural networks on graphs. In *IEEE conference on computer vision and pattern recognition*, pages 3693–3702, 2017.
- Uriel Singer, Ido Guy, and Kira Radinsky. Node embedding over temporal graphs. *arXiv preprint arXiv:1903.08889*, 2019.
- Richard Socher, Danqi Chen, Christopher D Manning, and Andrew Ng. Reasoning with neural tensor networks for knowledge base completion. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 926–934, 2013.
- Gustav Sourek, Vojtech Aschenbrenner, Filip Zelezny, Steven Schockaert, and Ondrej Kuzelka. Lifted relational neural networks: Efficient learning of latent relational structures. *Journal of Artificial Intelligence Research (JAIR)*, 62:69–100, 2018.
- Myra Spiliopoulou. Evolution in social networks: A survey. In *Social Network Data Analytics*, pages 149–175. Springer, 2011.
- Peter Stange. On the efficient update of the singular value decomposition. *PAMM*, 8:10827–10828, 2008.
- Yizhou Sun and Jiawei Han. Mining heterogeneous information networks: a structural analysis approach. *Acm Sigkdd Explorations Newsletter*, 14(2):20–28, 2013.
- Yizhou Sun, Jiawei Han, Charu C Aggarwal, and Nitesh V Chawla. When will it happen?: relationship prediction in heterogeneous information networks. In *Fifth ACM international Conference on Web Search and Data Mining*, pages 663–672. ACM, 2012.
- Daniel L. Sussman, Minh Tang, Donniell E. Fishkind, and Carey E. Priebe. A consistent adjacency spectral embedding for stochastic blockmodel graphs. *Journal of the American Statistical Association*, 107(499):1119–1128, 2012.
- Charles Sutton, Andrew McCallum, and Khashayar Rohanimanesh. Dynamic conditional random fields: Factorized probabilistic models for labeling and segmenting sequence data. *Journal of Machine Learning Research*, 8(Mar):693–723, 2007.
- Aynaz Taheri, Kevin Gimpel, and Tanya Berger-Wolf. Learning to represent the evolution of dynamic graphs with recurrent models. In *Companion Proceedings of The 2019 World Wide Web Conference, WWW '19*, pages 301–307, New York, NY, USA, 2019. ACM. ISBN 978-1-4503-6675-5. doi: 10.1145/3308560.3316581. URL <http://doi.acm.org/10.1145/3308560.3316581>.
- Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. Line: Large-scale information network embedding. In *Twenty-Fourth International Conference on World Wide Web*, pages 1067–1077. International World Wide Web Conferences Steering Committee, 2015.

- Lei Tang and Huan Liu. Relational learning via latent social dimensions. In *Fifteenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '09, pages 817–826, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-495-9. doi: 10.1145/1557019.1557109. URL <http://doi.acm.org/10.1145/1557019.1557109>.
- Phi Vu Tran. Learning to make predictions on graphs with autoencoders. In *2018 IEEE 5th International Conference on Data Science and Advanced Analytics (DSAA)*, pages 237–245. IEEE, 2018.
- Volker Tresp, Cristóbal Esteban, Yinchong Yang, Stephan Baier, and Denis Krompaß. Learning with memory embeddings. In *NIPS 2015 Workshop on Nonparametric Methods for Large Scale Representation Learning (extended TR)*, 2015.
- Volker Tresp, Yunpu Ma, Stephan Baier, and Yinchong Yang. Embedding learning for declarative memories. In *European Semantic Web Conference*, pages 202–216. Springer, 2017.
- Rakshit Trivedi, Hanjun Dai, Yichen Wang, and Le Song. Know-evolve: Deep temporal reasoning for dynamic knowledge graphs. In *International Conference on Machine Learning (ICML)*, pages 3462–3471. JMLR. org, 2017.
- Rakshit Trivedi, Mehrdad Farajtabar, Prasenjeet Biswal, and Hongyuan Zha. Dyrep: Learning representations over dynamic graphs. In *International Conference on Learning Representations (ICLR)*, 2019.
- Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. Complex embeddings for simple link prediction. In *International Conference on Machine Learning (ICML)*, pages 2071–2080, 2016.
- Théo Trouillon, Christopher R Dance, Éric Gaussier, Johannes Welbl, Sebastian Riedel, and Guillaume Bouchard. Knowledge graph completion via complex tensor factorization. *The Journal of Machine Learning Research*, 18(1):4735–4772, 2017.
- Guy Van den Broeck. On the completeness of first-order knowledge compilation for lifted probabilistic inference. In *Advances in Neural Information Processing Systems*, pages 1386–1394, 2011.
- Guy Van den Broeck and Adnan Darwiche. On the complexity and approximation of binary evidence in lifted inference. In *Advances in Neural Information Processing Systems*, pages 2868–2876, 2013.
- Jan Van Haaren, Guy Van den Broeck, Wannes Meert, and Jesse Davis. Lifted generative learning of markov logic networks. *Machine Learning*, 103(1):27–55, 2015.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 5998–6008, 2017.

- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. In *International Conference on Learning Representations (ICLR)*, 2018.
- Daixin Wang, Peng Cui, and Wenwu Zhu. Structural deep network embedding. In *Twenty-Second ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1225–1234. ACM, 2016.
- Minjie Wang and Quan Ga. Deep graph library, 2018.
- Quan Wang, Zhendong Mao, Bin Wang, and Li Guo. Knowledge graph embedding: A survey of approaches and applications. *IEEE Transactions on Knowledge and Data Engineering*, 29(12):2724–2743, 2017a.
- Tong Wang, Xing-Sheng He, Ming-Yang Zhou, and Zhong-Qian Fu. Link prediction in evolving networks based on popularity of nodes. *Scientific reports*, 2017b.
- Xiang Wang, Xiangnan He, Yixin Cao, Meng Liu, and Tat-Seng Chua. Kgat: Knowledge graph attention network for recommendation. *arXiv preprint arXiv:1905.07854*, 2019.
- Xiaolong Wang and Abhinav Gupta. Videos as space-time region graphs. In *European Conference on Computer Vision (ECCV)*, pages 399–417, 2018.
- Yanjie Wang, Rainer Gemulla, and Hui Li. On multi-relational link prediction with bilinear models. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Waloddi Weibull et al. A statistical distribution function of wide applicability. *Journal of Applied Mechanics*, 18(3):293–297, 1951.
- Boris Weisfeiler and Andrei A Lehman. A reduction of a graph to a canonical form and an algebra arising during this reduction. *Nauchno-Technicheskaya Informatsia*, 2(9):12–16, 1968.
- Greg Welch, Gary Bishop, et al. An introduction to the kalman filter. 1995.
- Jianfeng Wen, Jianxin Li, Yongyi Mao, Shini Chen, and Richong Zhang. On the representation and embedding of knowledge bases beyond binary relations. In *Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.
- Chao-Yuan Wu, Amr Ahmed, Alex Beutel, Alexander J Smola, and How Jing. Recurrent recommender networks. In *Tenth ACM international Conference on Web Search and Data Mining*, pages 495–503. ACM, 2017.
- Tianxing Wu, Arijit Khan, Huan Gao, and Cheng Li. Efficiently embedding dynamic knowledge graphs. *arXiv preprint arXiv:1910.06708*, 2019.
- Yu Xin, Zhi-Qiang Xie, and Jing Yang. An adaptive random walk sampling method on dynamic community detection. *Expert Systems with Applications*, 58:10–19, 2016.

- Liang Xiong, Xi Chen, Tzu-Kuo Huang, Jeff Schneider, and Jaime G Carbonell. Temporal collaborative filtering with bayesian probabilistic tensor factorization. In *SIAM International Conference on Data Mining*, pages 211–222. SIAM, 2010.
- Chengjin Xu, Mojtaba Nayyeri, Fouad Alkhoury, Jens Lehmann, and Hamed Shariat Yazdi. Temporal knowledge graph embedding model based on additive time series decomposition. *arXiv preprint arXiv:1911.07893*, 2019a.
- Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations (ICLR)*, 2019b.
- Naganand Yadati, Madhav Nimishakavi, Prateek Yadav, Anand Louis, and Partha Talukdar. Hypergnn: Hypergraph convolutional networks for semi-supervised classification. *arXiv preprint arXiv:1809.02589*, 2018.
- Pinar Yanardag and SVN Vishwanathan. A structural smoothing framework for robust graph comparison. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2134–2142, 2015.
- Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding entities and relations for learning and inference in knowledge bases. In *International Conference on Learning Representations (ICLR)*, 2015.
- Yang Yang, Nitesh Chawla, Yizhou Sun, and Jiawei Han. Predicting links in multi-relational and heterogeneous networks. In *2012 IEEE 12th international conference on data mining*, pages 755–764. IEEE, 2012.
- Lin Yao, Luning Wang, Lv Pan, and Kai Yao. Link prediction based on common-neighbors for dynamic social network. *Procedia Computer Science*, pages 82–89, 2016.
- Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec. Graph convolutional neural networks for web-scale recommender systems. In *Twenty-Fourth ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 974–983. ACM, 2018a.
- Zhitao Ying, Jiaxuan You, Christopher Morris, Xiang Ren, Will Hamilton, and Jure Leskovec. Hierarchical graph representation learning with differentiable pooling. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 4800–4810, 2018b.
- Bing Yu, Haoteng Yin, and Zhanxing Zhu. Spatio-temporal graph convolutional networks: a deep learning framework for traffic forecasting. In *Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3634–3640. AAAI Press, 2018a.
- Bing Yu, Mengzhang Li, Jiyong Zhang, and Zhanxing Zhu. 3d graph convolutional networks with temporal graphs: A spatial information free framework for traffic forecasting. *arXiv preprint arXiv:1903.00919*, 2019.
- Wenchao Yu, Wei Cheng, Charu C Aggarwal, Haifeng Chen, and Wei Wang. Link prediction with spatial and temporal consistency in dynamic networks. In *Twenty-Sixth International*

- Joint Conference on Artificial Intelligence (IJCAI)*, pages 3343–3349, 2017a. doi: 10.24963/ijcai.2017/467. URL <https://doi.org/10.24963/ijcai.2017/467>.
- Wenchao Yu, Wei Cheng, Charu C Aggarwal, Kai Zhang, Haifeng Chen, and Wei Wang. Netwalk: A flexible deep embedding approach for anomaly detection in dynamic networks. In *Twenty-Fourth ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2672–2681. ACM, 2018b.
- Zhiyong Yu, Jijie Chen, Kun Quo, Yuzhong Chen, and Qian Xu. Overlapping community detection based on random walk and seeds extension. In *Twelfth Chinese Conference on Computer Supported Cooperative Work and Social Computing*, pages 18–24. ACM, 2017b.
- Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan R Salakhutdinov, and Alexander J Smola. Deep sets. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 3391–3401, 2017.
- Ce Zhang. Deepdive: a data management system for automatic knowledge base construction. *University of Wisconsin-Madison, Madison, Wisconsin*, 2015.
- Daokun Zhang, Jie Yin, Xingquan Zhu, and Chengqi Zhang. Metagraph2vec: complex semantic path augmented heterogeneous network embedding. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 196–208. Springer, 2018a.
- Daokun Zhang, Jie Yin, Xingquan Zhu, and Chengqi Zhang. Network representation learning: A survey. *IEEE Transactions on Big Data*, 2018b.
- Jian Zhang. A survey on streaming algorithms for massive graphs. In *Managing and Mining Graph Data*, pages 393–420. Springer, 2010.
- Jiani Zhang, Xingjian Shi, Junyuan Xie, Hao Ma, Irwin King, and Dit-Yan Yeung. GaAN: Gated attention networks for learning on large and spatiotemporal graphs. In *Uncertainty in Artificial Intelligence*, 2018c.
- Ziwei Zhang, Peng Cui, Jian Pei, Xiao Wang, and Wenwu Zhu. Timers: Error-bounded svd restart on dynamic networks. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018d.
- Jie Zhou, Ganqu Cui, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, and Maosong Sun. Graph neural networks: A review of methods and applications. *arXiv preprint arXiv:1812.08434*, 2018a.
- Lekui Zhou, Yang Yang, Xiang Ren, Fei Wu, and Yueting Zhuang. Dynamic network embedding by modeling triadic closure process. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018b.
- Jia Zhu, Qing Xie, and Eun Jung Chin. A hybrid time-series link prediction framework for large social network. In *International Conference on Database and Expert Systems Applications*, pages 345–359. Springer, 2012.

- Linhong Zhu, Dong Guo, Junming Yin, Greg Ver Steeg, and Aram Galstyan. Scalable temporal latent space inference for link prediction in dynamic social networks. *IEEE Transactions on Knowledge and Data Engineering*, 28(10):2765–2777, 2016.
- Yu Zhu, Hao Li, Yikang Liao, Beidou Wang, Ziyu Guan, Haifeng Liu, and Deng Cai. What to do next: Modeling user behaviors by time-LSTM. In *Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3602–3608, 2017.
- Yuan Zuo, Guannan Liu, Hao Lin, Jia Guo, Xiaoqian Hu, and Junjie Wu. Embedding temporal network via neighborhood formation. In *Twenty-Fourth ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2857–2866. ACM, 2018.