

Dynamic Graph Convolutional Networks Using the Tensor M-Product

Osman Asif Malik* Shashanka Ubaru† Lior Horesh‡ Misha E. Kilmer§
Haim Avron¶

Abstract

Many irregular domains such as social networks, financial transactions, neuron connections, and natural language constructs are represented using graph structures. In recent years, a variety of graph neural networks (GNNs) have been successfully applied for representation learning and prediction on such graphs. In many of the real-world applications, the underlying graph changes over time, however, most of the existing GNNs are inadequate for handling such dynamic graphs. In this paper we propose a novel technique for learning embeddings of dynamic graphs using a tensor algebra framework. Our method extends the popular graph convolutional network (GCN) for learning representations of dynamic graphs using the recently proposed tensor M-product technique. Theoretical results presented establish a connection between the proposed tensor approach and spectral convolution of tensors. The proposed method TM-GCN is consistent with the Message Passing Neural Network (MPNN) framework, accounting for both spatial and temporal message passing. Numerical experiments on real-world datasets demonstrate the performance of the proposed method for edge classification and link prediction tasks on dynamic graphs. We also consider an application related to the COVID-19 pandemic, and show how our method can be used for early detection of infected individuals from contact tracing data.

1 Introduction

Graphs are popular data structures used to effectively represent interactions and structural relationships between entities in structured data domains. Inspired by the success of deep neural networks for learning representations in the image and language domains, recently, application of neural networks for graph representation learning has attracted much interest. A number of graph neural network (GNN) architectures have been explored in the contemporary literature for a variety of graph related tasks and applications [33, 30]. Methods based




Figure 1: Our proposed TM-GCN approach.

on graph convolution filters which extend convolutional neural networks (CNNs) to irregular graph domains are popular [4, 7, 12]. Most of these GNN models operate on a given, static graph.

In many real-world applications, the underlying graph changes over time, and learning representations of such dynamic graphs is essential. Examples include analyzing social networks [2], detecting fraud and crime in financial networks [23], traffic control [32], understanding neuronal activities in the brain [6], and analyzing contact tracing data [28]. In such dynamic settings, the temporal interdependence in the graph connections and features also play a substantial role. However, efficient GNN methods that handle time varying graphs and that capture the temporal correlations are lacking.

By *dynamic graph*, we refer to a sequence of graphs $\mathcal{G}^{(t)} = (V, \mathbf{A}^{(t)}, \mathbf{X}^{(t)})$, $t \in \{1, 2, \dots, T\}$, with a fixed set V of N nodes, adjacency matrices $\mathbf{A}^{(t)} \in \mathbb{R}^{N \times N}$, and graph feature matrices $\mathbf{X}^{(t)} \in \mathbb{R}^{N \times F}$ where $\mathbf{X}_n^{(t)} \in \mathbb{R}^F$ is the feature vector consisting of F features associated with node n at time t . The graphs can be weighted, and directed or undirected. They can also have additional properties like (time varying) node and edge classes, which would be stored in a separate structure. Suppose we only observe the first $T' < T$ graphs in the sequence. The goal of our method is to use these observations to predict some property of the remaining $T - T'$ graphs. In this paper, we consider edge classification, link prediction and node property prediction tasks.

In recent years, tensor constructs have been explored to effectively process high-dimensional data, in order

*University of Colorado Boulder, osman.malik@colorado.edu

†IBM Research, shashanka.ubaru@ibm.com

‡IBM Research, lhoresh@us.ibm.com

§Tufts University, misha.kilmer@tufts.edu

¶Tel Aviv University, haimav@tauex.tau.ac.il

to better leverage the multidimensional structure of such data [13]. Tensor based approaches have been shown to perform well in many applications. Recently, a new tensor framework called the *tensor M-product framework* [3, 10] was proposed that extends matrix based theory to high-dimensional architectures.

In this paper, we propose a novel tensor variant of the popular graph convolutional network (GCN) architecture [12], which we call TM-GCN. It captures correlation over time by leveraging the tensor M-product framework. The flexibility and matrix mimeticability of the framework, help us adapt the GCN architecture to tensor space. Figure 1 illustrates our method at a high level: First, the time varying adjacency matrices $\mathbf{A}^{(t)}$ and feature matrices $\mathbf{X}^{(t)}$ of the dynamic graph are aggregated into an adjacency tensor and a feature tensor, respectively. These tensors are then fed into our TM-GCN, which computes an embedding that can be used for a variety of tasks, such as link prediction, and edge and node classification. GCN architectures are motivated by graph convolution filtering, i.e., applying filters/functions to the graph Laplacian [4], and we establish a similar connection between TM-GCN and spectral filtering of tensors. Such results suggest possible extensions of other convolution based GNNs such as [4, 7] for dynamic graphs using the tensor framework. The Message Passing Neural Network (MPNN) framework has been used to describe spatial convolution GNNs [8]. We show that TM-GCN is consistent with the MPNN framework, and accounts for spatial and temporal message passing. Experimental results on real datasets illustrate the performance of our method for the edge classification and link prediction tasks on dynamic graphs. We also demonstrate how TM-GCN can be used in an important application related to the COVID-19 pandemic. We show how GNNs can be used for early identification of individuals who are infected (potentially before they display symptoms) from contact tracing data and a dynamic graph based SEIR model [28].

2 Related Work

Unsupervised Embedding: Unsupervised graph embedding techniques have been popular for link prediction on static graphs [5]. A number of dynamic graph embedding methods have been proposed recently, which extend the static ones. DANE [16] adapted the popular dimensionality reduction approaches such as Eigenmaps to time varying graphs by efficiently updating the eigenvectors from the prior ones. The popular random walk based methods have also been extended to obey the temporal order in recent works [22].

Numerous deep neural network based unsupervised learning methods have been developed for dynamic

graph embedding. Examples include DynGEM [9], Know-Evolve [26], DyRep [27], Dynamic-Triad [34], and others. In most of these methods, a temporal smoothness regularization is used to obtain stable embedding across consecutive time-steps.

Supervised Learning: The idea of using graph convolution based on the spectral graph theory for GNNs was first introduced by [4]. [7] then proposed *Chebnet*, where the spectral filter was approximated by Chebyshev polynomials in order to make it faster and localized. [12] presented the simplified GCN, a degree-one polynomial approximation of Chebnet, in order to speed up computation further and improve the performance. There are many other works that deal with GNNs when the graph and features are fixed/static; see the review papers [33] and [30] and references therein.

Recently, Li et. al [17] develop a diffusion convolutional RNN for traffic forecasting, where road networks are modeled assuming both the nodes and edges remain fixed over time, unlike in our setting. Seo et. al [25] devise the Graph Convolutional Recurrent Network for graphs with time varying features, while the edges are fixed over time. EdgeConv was proposed in [29], which is a neural network (NN) approach that applies convolution operations on static graphs in a dynamic fashion. [32] develop a temporal GCN method called T-GCN, which they apply for traffic prediction. Here too, the graph remains fixed over time, and only the features vary. [31] propose a method which they refer to as a tensor graph CNN. Here, the standard GCN [12] based on matrix algebra is considered, and a “cross graph convolution” layer is introduced to handle the time varying aspect of the dynamic graph. In particular, the cross graph convolution layer involves computing a parameterized Kronecker sum of the current adjacency matrix with the previously processed adjacency matrix, followed by a GCN layer. Recently, [18] described a tensor version of GCN for text classification, where the text semantics are represented as a three-dimensional graph tensor. This work neither considers time varying graphs, nor the tensor M-product framework.

The set of methods most relevant to our setting of learning embeddings of dynamic graphs use combinations of GNNs and recurrent architectures (RNN), to capture the graph structure and handle time dynamics, respectively. The approach in [19] uses Long Short-Term Memory (LSTM), a recurrent network, in order to handle time variations along with GNNs. They design architectures for semi-supervised node classification and for supervised graph classification. [23] presented a variant of GCN called *EvolveGCN*, where Gated Recurrent Units (GRUs) and LSTMs are coupled with a GCN to handle dynamic graphs. This paper is currently the

state-of-the-art. [24] proposed the use of a temporal self-attention layer for dynamic graph representation learning. However, all these approaches are based on a heuristic RNN/GRU mechanism to evolve weights, and the models are not time aware (time is not an explicit entity). [21] present a tensor NN which utilizes the tensor M-product framework. Their approach is applicable to image and other high-dimensional data that lie on regular grids.

3 Tensor M-Product Framework

Here, we cover the necessary preliminaries on tensors and the M-product framework. For a more general introduction to tensors, we refer the reader to the review paper [13]. In the present paper, a *tensor* is a three-dimensional array of real numbers denoted by boldface Euler script letters, e.g. $\mathcal{X} \in \mathbb{R}^{I \times J \times T}$. Matrices are denoted by bold uppercase letters, e.g. \mathbf{X} ; vectors are denoted by bold lowercase letter, e.g. \mathbf{x} ; and scalars are denoted by lowercase letters, e.g. x . An element at position (i, j, t) in a tensor is denoted by subscripts, e.g. \mathcal{X}_{ijt} , with similar notation for elements of matrices and vectors. A colon will denote all elements along that dimension; $\mathbf{X}_{:i}$ denotes the i th row of the matrix \mathbf{X} , and $\mathcal{X}_{::k}$ denotes the k th frontal slice of \mathcal{X} . The vectors $\mathcal{X}_{ij:}$ are called the *tubes* of \mathcal{X} .

The framework we consider relies on a new definition of the product of two tensors, called the M-product [3, 11, 10]. A distinguishing feature of this framework is that the M-product of two three-dimensional tensors is also three-dimensional, which is not the case for e.g. tensor contractions [13]. It allows one to elegantly generalize many classical numerical methods from linear algebra. The framework, originally developed for three-dimensional tensors, has been extended to handle tensors of dimension greater than three [11]. The following definitions 3.1–3.3 describe the M-product.

DEFINITION 3.1. (M-TRANSFORM) Let $\mathbf{M} \in \mathbb{R}^{T \times T}$ be a mixing matrix. The M-transform of a tensor $\mathcal{X} \in \mathbb{R}^{I \times J \times T}$ is denoted by $\mathcal{X} \times_3 \mathbf{M} \in \mathbb{R}^{I \times J \times T}$ and defined elementwise as

$$(3.1) \quad (\mathcal{X} \times_3 \mathbf{M})_{ijt} \stackrel{\text{def}}{=} \sum_{k=1}^T \mathbf{M}_{tk} \mathcal{X}_{ijk}.$$

We say that $\mathcal{X} \times_3 \mathbf{M}$ is in the transformed space. Note that if \mathbf{M} is invertible, then $(\mathcal{X} \times_3 \mathbf{M}) \times_3 \mathbf{M}^{-1} = \mathcal{X}$. Consequently, $\mathcal{X} \times_3 \mathbf{M}^{-1}$ is the inverse M-transform of \mathcal{X} . The definition in (3.1) may also be written in matrix form as $\mathcal{X} \times_3 \mathbf{M} \stackrel{\text{def}}{=} \text{fold}(\mathbf{M} \text{unfold}(\mathcal{X}))$, where the *unfold* operation takes the tubes of \mathcal{X} and stack them as columns into a $T \times IJ$ matrix, and $\text{fold}(\text{unfold}(\mathcal{X})) = \mathcal{X}$.

DEFINITION 3.2. (FACEWISE PRODUCT) Let $\mathcal{X} \in \mathbb{R}^{I \times J \times T}$ and $\mathcal{Y} \in \mathbb{R}^{J \times K \times T}$ be two tensors. The facewise product, denoted by $\mathcal{X} \triangle \mathcal{Y} \in \mathbb{R}^{I \times K \times T}$, is defined facewise as $(\mathcal{X} \triangle \mathcal{Y})_{::t} \stackrel{\text{def}}{=} \mathcal{X}_{::t} \mathcal{Y}_{::t}$.

DEFINITION 3.3. (M-PRODUCT) Let $\mathcal{X} \in \mathbb{R}^{I \times J \times T}$ and $\mathcal{Y} \in \mathbb{R}^{J \times K \times T}$ be two tensors, and let $\mathbf{M} \in \mathbb{R}^{T \times T}$ be an invertible matrix. The M-product, denoted by $\mathcal{X} \star \mathcal{Y} \in \mathbb{R}^{I \times K \times T}$, is defined as

$$\mathcal{X} \star \mathcal{Y} \stackrel{\text{def}}{=} ((\mathcal{X} \times_3 \mathbf{M}) \triangle (\mathcal{Y} \times_3 \mathbf{M})) \times_3 \mathbf{M}^{-1}.$$

In the original formulation of the M-product, \mathbf{M} was chosen to be the Discrete Fourier Transform (DFT) matrix, which allows efficient computation using the Fast Fourier Transform (FFT) [3, 11]. The framework was later extended for arbitrary invertible \mathbf{M} (e.g. discrete cosine and wavelet transforms) [10]. Additional details are in the supplement.

4 Tensor Dynamic Graph Embedding

Our approach is inspired by the first order GCN by [12] for static graphs, owed to its simplicity and effectiveness. For a graph with adjacency matrix \mathbf{A} and feature matrix \mathbf{X} , a GCN layer takes the form $\mathbf{Y} = \sigma(\tilde{\mathbf{A}}\mathbf{X}\mathbf{W})$, where

$$\tilde{\mathbf{A}} \stackrel{\text{def}}{=} \tilde{\mathbf{D}}^{-1/2}(\mathbf{A} + \mathbf{I})\tilde{\mathbf{D}}^{-1/2},$$

$\tilde{\mathbf{D}}$ is diagonal with $\tilde{\mathbf{D}}_{ii} = 1 + \sum_j \mathbf{A}_{ij}$, \mathbf{I} is the matrix identity, \mathbf{W} is a matrix to be learned when training the NN, and σ is an activation function, e.g., ReLU. Our approach translates this to a tensor model by utilizing the M-product framework. We first introduce a tensor activation function $\hat{\sigma}$ which operates in the transformed space.

DEFINITION 4.1. Let $\mathcal{A} \in \mathbb{R}^{I \times J \times T}$ be a tensor and σ an elementwise activation function. We define the activation function $\hat{\sigma}$ as $\hat{\sigma}(\mathcal{A}) \stackrel{\text{def}}{=} \sigma(\mathcal{A} \times_3 \mathbf{M}) \times_3 \mathbf{M}^{-1}$.

We can now define our proposed dynamic graph embedding. Let $\mathcal{A} \in \mathbb{R}^{N \times N \times T}$ be a tensor with frontal slices $\mathcal{A}_{::t} = \tilde{\mathbf{A}}^{(t)}$, where $\tilde{\mathbf{A}}^{(t)}$ is the normalization of $\mathbf{A}^{(t)}$. Moreover, let $\mathcal{X} \in \mathbb{R}^{N \times F \times T}$ be a tensor with frontal slices $\mathcal{X}_{::t} = \mathbf{X}^{(t)}$. Finally, let $\mathcal{W} \in \mathbb{R}^{F \times F' \times T}$ be a weight tensor. We define our dynamic graph embedding as $\mathcal{Y} = \mathcal{A} \star \hat{\sigma}(\mathcal{A} \star \mathcal{X} \star \mathcal{W}^{(0)}) \star \mathcal{W}^{(1)}$. This computation can also be repeated in multiple layers. For example, a 2-layer formulation would be of the form

$$\mathcal{Y} = \mathcal{A} \star \hat{\sigma}(\mathcal{A} \star \mathcal{X} \star \mathcal{W}^{(0)}) \star \mathcal{W}^{(1)}.$$

One important consideration is how to choose the matrix \mathbf{M} which defines the M-product. For time-varying graphs, we choose \mathbf{M} to be lower triangular

and banded so that each frontal slice $(\mathcal{A} \times_3 \mathbf{M})_{::t}$ is a linear combination of the adjacency matrices $\mathcal{A}_{::\max(1,t-b+1)}, \dots, \mathcal{A}_{::t}$, where we refer to b as the “bandwidth” of \mathbf{M} . This choice ensures that each frontal slice $(\mathcal{A} \times_3 \mathbf{M})_{::t}$ only contains information from current and past graphs that are close temporally. We consider two variants of the lower banded triangular \mathbf{M} matrix in the experiments; see the supplement for details. Another possibility is to treat \mathbf{M} as a parameter matrix to be learned from the data.

In order to avoid over-parameterization and improve the performance, we choose the weight tensor \mathcal{W} (at each layer), such that each of the frontal slices of \mathcal{W} in the transformed domain remains the same, i.e., $(\mathcal{W} \times_3 \mathbf{M})_{::t} = (\mathcal{W} \times_3 \mathbf{M})_{::t'} \forall t, t'$. In other words, the parameters in each layer are shared and learned over all the training instances. This reduces the number of parameters to be learned significantly.

An embedding $\mathbf{Y} \in \mathbb{R}^{N \times F' \times T}$ can now be used for various prediction tasks, like link prediction, and edge and node classification. In Section 5, we apply our method for edge classification and link prediction by using a model similar to that used by [23]: Given an edge between nodes m and n at time t , the predictive model is

$$(4.2) \quad p(m, n, t) \stackrel{\text{def}}{=} \text{softmax}(\mathbf{U}[(\mathbf{Y} \times_3 \mathbf{M})_{m:t}, (\mathbf{Y} \times_3 \mathbf{M})_{n:t}]^\top),$$

where $(\mathbf{Y} \times_3 \mathbf{M})_{m:t} \in \mathbb{R}^{F'}$ and $(\mathbf{Y} \times_3 \mathbf{M})_{n:t} \in \mathbb{R}^{F'}$ are row vectors, $\mathbf{U} \in \mathbb{R}^{C \times 2F'}$ is a weight matrix, and C the number of classes. Note that the embedding \mathbf{Y} is first M-transformed before the matrix \mathbf{U} is applied to the appropriate feature vectors. This, combined with the fact that the tensor activation functions are applied elementwise in the transformed domain, allow us to avoid ever needing to apply the inverse M-transform. This approach reduces the computational cost, and has been found to improve performance in the edge classification task.

4.1 Theoretical Motivation for TM-GCN Here, we present the results that establish the connection between the proposed TM-GCN and spectral convolution of tensors, in particular spectral filtering and approximation on dynamic graphs. This is analogous to the graph convolution based on spectral graph theory in the GNNs by [4], [7], and [12]. All proofs and additional details are provided in Section D of the supplement.

Let $\mathcal{L} \in \mathbb{R}^{N \times N \times T}$ be a form of tensor Laplacian defined as $\mathcal{L} \stackrel{\text{def}}{=} \mathcal{J} - \mathcal{A}$. Throughout the remainder of this subsection, we will assume that the adjacency matrices $\mathbf{A}^{(t)}$ are symmetric.

PROPOSITION 4.1. *The tensor \mathcal{L} has an eigendecompo-*

sition $\mathcal{L} = \mathcal{Q} \star \mathcal{D} \star \mathcal{Q}^\top$.

DEFINITION 4.2. (FILTERING) *Given a signal $\mathcal{X} \in \mathbb{R}^{N \times 1 \times T}$ and a function $g : \mathbb{R}^{1 \times 1 \times T} \rightarrow \mathbb{R}^{1 \times 1 \times T}$, we define the tensor spectral graph filtering of \mathcal{X} with respect to g as*

$$(4.3) \quad \mathcal{X}_{\text{filt}} \stackrel{\text{def}}{=} \mathcal{Q} \star g(\mathcal{D}) \star \mathcal{Q}^\top \star \mathcal{X},$$

where

$$g(\mathcal{D})_{mn} \stackrel{\text{def}}{=} \begin{cases} g(\mathcal{D}_{mn}) & \text{if } m = n, \\ \mathbf{0} & \text{if } m \neq n. \end{cases}$$

In order to avoid the computation of an eigendecomposition, [7] uses a polynomial to approximate the filter function. We take a similar approach, and approximate $g(\mathcal{D})$ with an M-product polynomial. For this approximation, we impose additional structure on g .

ASSUMPTION 1. *Assume that $g : \mathbb{R}^{1 \times 1 \times T} \rightarrow \mathbb{R}^{1 \times 1 \times T}$ is defined as*

$$g(\mathcal{V}) \stackrel{\text{def}}{=} f(\mathcal{V} \times_3 \mathbf{M}) \times_3 \mathbf{M}^{-1},$$

where f is defined elementwise as $f(\mathcal{V} \times_3 \mathbf{M})_{11t} \stackrel{\text{def}}{=} f^{(t)}((\mathcal{V} \times_3 \mathbf{M})_{11t})$ with each $f^{(t)} : \mathbb{R} \rightarrow \mathbb{R}$ continuous.

PROPOSITION 4.2. *Suppose g satisfies Assumption 1. For any $\varepsilon > 0$, there exists an integer K and a set $\{\theta^{(k)}\}_{k=1}^K \subset \mathbb{R}^{1 \times 1 \times T}$ such that*

$$\left\| g(\mathcal{D}) - \sum_{k=0}^K \mathcal{D}^{\star k} \star \theta^{(k)} \right\| < \varepsilon,$$

where $\|\cdot\|$ is the tensor Frobenius norm, and where $\mathcal{D}^{\star k} \stackrel{\text{def}}{=} \mathcal{D} \star \dots \star \mathcal{D}$ is the M-product of k instances of \mathcal{D} , with the convention that $\mathcal{D}^{\star 0} = \mathcal{J}$.

As in the work of [7], a tensor polynomial approximation allows us to approximate $\mathcal{X}_{\text{filt}}$ in (4.3) without computing the eigendecomposition of \mathcal{L} :

$$(4.4) \quad \begin{aligned} \mathcal{X}_{\text{filt}} &= \mathcal{Q} \star g(\mathcal{D}) \star \mathcal{Q}^\top \star \mathcal{X} \\ &\approx \mathcal{Q} \star \left(\sum_{k=0}^K \mathcal{D}^{\star k} \star \theta^{(k)} \right) \star \mathcal{Q}^\top \star \mathcal{X} \\ &= \left(\sum_{k=0}^K \mathcal{L}^{\star k} \star \theta^{(k)} \right) \star \mathcal{X}. \end{aligned}$$

All that is necessary is to compute tensor powers of \mathcal{L} . We can also define tensor polynomial analogs of the Chebyshev polynomials and do the approximation in (4.4) in terms of those instead of the tensor monomials

\mathcal{D}^{*k} . We note that if a degree-one approximation is used, the computation in (4.4) becomes

$$\begin{aligned}\mathcal{X}_{\text{filt}} &\approx (\mathcal{J} \star \boldsymbol{\theta}^{(0)} + \mathcal{L} \star \boldsymbol{\theta}^{(1)}) \star \mathcal{X} \\ &= (\mathcal{J} \star \boldsymbol{\theta}^{(0)} + (\mathcal{J} - \mathcal{A}) \star \boldsymbol{\theta}^{(1)}) \star \mathcal{X}.\end{aligned}$$

Setting $\boldsymbol{\theta} \stackrel{\text{def}}{=} \boldsymbol{\theta}^{(0)} = -\boldsymbol{\theta}^{(1)}$, which is analogous to the parameter choice made in the degree-one approximation in [12], we get

$$(4.5) \quad \mathcal{X}_{\text{filt}} \approx \mathcal{A} \star \mathcal{X} \star \boldsymbol{\theta}.$$

If we let \mathcal{X} contain F signals, i.e., $\mathcal{X} \in \mathbb{R}^{N \times F \times T}$, and apply F' filters, (4.5) becomes

$$\mathcal{X}_{\text{filt}} \approx \mathcal{A} \star \mathcal{X} \star \boldsymbol{\Theta} \in \mathbb{R}^{N \times F' \times T},$$

where $\boldsymbol{\Theta} \in \mathbb{R}^{F \times F' \times T}$. This is precisely our embedding model, with $\boldsymbol{\Theta}$ replaced by a learnable parameter tensor \mathcal{W} . These results show: (a) the connection between TM-GCN and spectral convolution of tensors, analogous to the GCN, and (b) that we can indeed develop higher order convolutional GNNs like [4, 7] for dynamic graphs using our framework.

4.2 Message Passing Framework The Message Passing Neural Network (MPNN) framework is popularly used to describe spatial convolution GNNs [8]. The graph convolution operation is considered to be a message passing process, with information being passed from one node to another along the edges. The message passing phase of MPNN constitutes updating the hidden state $\mathbf{h}_{v,\ell}$ at node v in the ℓ th layer with message $\mathbf{m}_{v,\ell+1}$ as

$$\begin{aligned}\mathbf{m}_{v,\ell+1} &= \sum_{w \in N(v)} \Phi_{\ell}(\mathbf{h}_{v,\ell}, \mathbf{h}_{w,\ell}, e_{vw}), \\ \mathbf{h}_{v,\ell+1} &= \Psi_{\ell}(\mathbf{h}_{v,\ell}, \mathbf{m}_{v,\ell+1}),\end{aligned}$$

where $N(v)$ is the neighbors of v in the graph, e_{vw} is the edge between nodes v and w , Φ_{ℓ} is a message function, and Ψ_{ℓ} is an update function. A number of GNN models can be defined using this standard MPNN framework for static graphs. For the standard GCN model [12], we have $\Phi_{\ell}(\mathbf{h}_{v,\ell}, \mathbf{h}_{w,\ell}) = A_{v,w} \mathbf{h}_{w,\ell}$, where $A_{v,w}$ is the entry of adjacency matrix \mathcal{A} , and $\Psi_{\ell}(\mathbf{h}_{v,\ell}, \mathbf{m}_{v,\ell+1}) = \sigma(\mathbf{m}_{v,\ell+1})$ where σ is a pointwise non-linear function, e.g., ReLU.

In this paper, we consider dynamic graphs and the designed GNN has to do spatial and temporal message passing. That is, for a graph $\mathcal{G}^{(t)}$ at time t , the MPNN should be modeled such that the information/message is passed between neighboring nodes, as well as the corresponding nodes in the graphs $\{\mathcal{G}^{(t-1)}, \mathcal{G}^{(t-2)}, \dots, \mathcal{G}^{(1)}\}$. Recently, a spatio-temporal message passing framework was defined for video processing in computer vision [20].

Table 1: Dataset statistics.

Dataset	Nodes	Edges	T	Window		Partitioning		
				(days)	C	S_{train}	S_{val}	S_{test}
SBM	1,000	1,601,999	50	–	–	35	5	10
BitcoinOTC	6,005	35,569	135	14	2	95	20	20
BitcoinAlpha	7,604	24,173	135	14	2	95	20	20
Reddit	3,818	163,008	86	14	2	66	10	10
Chess	7,301	64,958	100	31	3	80	10	10

However, their framework does not account for time direction, and the graph is not considered to be evolving. We define the message passing framework for a dynamic graph as follows: The message passing phase will constitute updating the hidden state $\mathbf{h}_{v,\ell}^{(t)}$ at node v of graph $\mathcal{G}^{(t)}$ in the ℓ th layer with message $\mathbf{m}_{v,\ell+1}^{(t)}$ as

$$\begin{aligned}\mathbf{m}_{v,\ell+1}^{(t)} &= \sum_{w \in N(v)} \sum_{\tau=1}^t \Phi_{\ell}^{(\tau)} \left(\sum_{\tau=1}^t \Gamma_{\ell}^{(\tau)}(\mathbf{h}_{v,\ell}^{(\tau)}, \mathbf{h}_{w,\ell}^{(\tau)}, e_{vw}^{\tau}) \right), \\ \mathbf{h}_{v,\ell+1}^{(t)} &= \Psi_{\ell}(\mathbf{h}_{v,\ell}^{(t)}, \mathbf{m}_{v,\ell+1}^{(t)}).\end{aligned}$$

Here, the function $\Gamma_{\ell}^{(\tau)}$ accounts for the message passing between hidden states over different time $\tau \leq t$, and function $\phi_{\ell}^{(\tau)}$ accounts for message passing between neighbors $N(v)$ over time $\tau \leq t$. The model accounts for extensive spatio-temporal message passing.

For the proposed TM-GCN model, we have a function $\Gamma_{\ell}^{(\tau)}(\mathbf{h}_{v,\ell}^{(\tau)}) = M_{t,\tau} \mathbf{h}_{v,\ell}^{(\tau)}$, where $M_{t,\tau}$ is the (t, τ) entry of the mixing matrix \mathbf{M} . The message function is $\Phi_{\ell}^{(\tau)}(\mathbf{z}_{v,\ell}^{(t)}, \mathbf{h}_{w,\ell}^{(\tau)}) = M_{t,\tau} A_{v,w,\tau}$, where $\mathbf{z}_{v,\ell}^{(t)} = \sum_{\tau=1}^t \Gamma^{(\tau)}(\mathbf{h}_{v,\ell}^{(\tau)})$, and $A_{v,w,\tau}$ is the entry of the adjacency tensor \mathcal{A} . The update function is $\Psi_{\ell}(\mathbf{h}_{v,\ell}^{(t)}, \mathbf{m}_{v,\ell+1}^{(t)}) = \sigma(\mathbf{m}_{v,\ell+1}^{(t)})$ with an elementwise non-linear function σ .

Note that the above message passing model does not include the inverse transform \mathbf{M}^{-1} as in the definition of the M-products. This is because, the M-transform is responsible for the temporal message passing and undoing it is not necessary. In our experiments too, we found that transforming back (applying the inverse transform \mathbf{M}^{-1}) did not yield improved results as suggested by the above MPNN model. This does not affect any of the theory presented in the previous section since the spectral filtering is performed in the transformed domain (see the supplement for details).

5 Numerical Experiments

We first present results for edge classification and link prediction. We then show how we can use GNNs for predicting the state of individuals from COVID-19 contact tracing graphs.

Table 2: Results for edge classification. Performance measures are F1 score[†] or accuracy*. A higher value is better.

Method	Dataset			
	Bitcoin OTC [†]	Bitcoin Alpha [†]	Reddit [†]	Chess*
WD-GCN	0.3562	0.2533	0.2337	0.4311
EvolveGCN	0.3483	0.2273	0.2012	0.4351
GCN	0.3402	0.2381	0.1968	0.4342
TM-GCN - M1	0.3660	0.3243	0.2057	0.4708
TM-GCN - M2	0.4361	0.2466	0.1833	0.4513

Table 3: Results for link prediction. Performance measure is MAP. A higher value is better.

Method	Dataset				
	SBM	Bitcoin OTC	Bitcoin Alpha	Reddit	Chess
WD-GCN	0.9436	0.8071	0.8795	0.3896	0.1279
EvolveGCN	0.7620	0.6985	0.7722	0.2866	0.0915
GCN	0.9201	0.6847	0.7655	0.3099	0.0899
TM-GCN - M1	0.9684	0.8026	0.9318	0.2270	0.1882
TM-GCN - M2	0.9799	0.8458	0.9631	0.1405	0.1514




Figure 2: Partitioning of \mathcal{A} into training, validation and testing data.

5.1 Datasets and Preprocessing We consider five datasets (links to the datasets are in the supplement): The Bitcoin Alpha and OTC transaction datasets [23], the Reddit body hyperlink dataset [14], a chess results dataset [15], and SBM is the structure block matrix by [9]. The bitcoin datasets consist of transaction histories for users on two different platforms. Each node is a user, and each directed edge indicates a transaction and is labeled with an integer between -10 and 10 which indicates the senders trust for the receiver. We convert these labels to two classes: positive (trustworthy) and negative (untrustworthy). The Reddit dataset is built from hyperlinks from one subreddit to another. Each node represents a subreddit, and each directed edge is an interaction which is labeled with -1 for a hostile interaction or $+1$ for a friendly interaction. We only consider those subreddits which have a total of 20 interactions or more. In the chess dataset, each node is a player, and each directed edge represents a match with the source node being the white player and the target node being the black player. Each edge is labeled -1 for a black victory, 0 for a draw, and $+1$ for a

white victory. Table 1 summarizes the statistics for the different datasets, where T is total # graphs and C is the # classes. The SBM dataset has no labels and hence we use it only for link prediction.

The data is temporally partitioned into T graphs, with each graph containing data from a particular time window. Both T and the time window length can vary between datasets. For each node-time pair (n, t) in these graphs, we compute the number of outgoing and incoming edges and use these two numbers as features. The adjacency tensor \mathcal{A} is then constructed as described in Section 4. The T frontal slices of \mathcal{A} are divided into S_{train} training slices, S_{val} validation slices, and S_{test} testing slices, which come sequentially after each other; see Figure 2 and Table 1.

Since the adjacency matrices corresponding to graphs are very sparse for these datasets, we apply the same technique as [23] and add the entries of each frontal slice $\mathcal{A}_{::t}$ to the following $l - 1$ frontal slices $\mathcal{A}_{::t}, \dots, \mathcal{A}_{::(t+l-1)}$, where we refer to l as the “edge life.” Note that this only affects \mathcal{A} , and that the added edges are not treated as real edges in the classification and prediction problems.

The bitcoin and Reddit datasets are heavily skewed, with about 90% of edges labeled positively, and the remaining labeled negatively. Since the negative instances are more interesting to identify (e.g. to prevent financial fraud or online hostility), we use the F1 score to evaluate the edge classification experiments on these datasets, treating the negative edges as the ones we want to identify. The classes are more well-balanced in the

chess dataset, so we use accuracy to evaluate those edge classification experiments.

5.2 Graph Tasks For the link prediction experiments, we follow [23] and use negative sampling to construct non-existing edges, and use mean average precision (MAP) as a performance measure. The negative sampling is done so that 5% of edges are existing edges for each time slice, and all other edges are non-existing. Precise definitions of the different performance measures we use are given in Section B of the supplement.

For edge classification, we use an embedding $\mathbf{y}_{\text{train}} = \mathcal{A}_{::(1:S_{\text{train}})} \star \mathcal{X}_{::(1:S_{\text{train}})} \star \mathcal{W}$ for training. When computing the embeddings for the validation and testing data, we still need S_{train} frontal slices of \mathcal{A} , which we get by using a sliding window of slices. This is illustrated in Figure 2, where the green, blue and red blocks show the frontal slices used when computing the embeddings for the training, validation and testing data, respectively. The embeddings for the validation and testing data are $\mathbf{y}_{\text{val}} = \mathcal{A}_{::(S_{\text{val}}+1:S_{\text{train}}+S_{\text{val}})} \star \mathcal{X}_{::(S_{\text{val}}+1:S_{\text{train}}+S_{\text{val}})} \star \mathcal{W}$ and $\mathbf{y}_{\text{test}} = \mathcal{A}_{::(S_{\text{val}}+S_{\text{test}}+1:T)} \star \mathcal{X}_{::(S_{\text{val}}+S_{\text{test}}+1:T)} \star \mathcal{W}$, respectively. For link prediction, we use the same embeddings, with the only difference that the embedding blocks contain $S_{\text{train}} - 1$ slices. This is necessary since we want to use information up to time t to predict edge existence at time $t + 1$. Preliminary experiments with 2-layer architectures did not show convincing improvements in performance. We believe this is due to the fact that the datasets only have two features, and that a 1-layer architecture therefore is sufficient for extracting relevant information in the data.

For training, we use the cross entropy loss function:

$$\text{loss} = - \sum_t \sum_{(m,n) \in E_t} \sum_{c=1}^C \alpha_c f(m,n,t)_c \log(p(m,n,t)_c),$$

where $\alpha \in \mathbb{R}^C$ is a vector summing to 1 which contains the weight of each class. For edge classification, $f(m,n,t) \in \mathbb{R}^C$ is a one-hot vector encoding the true class of the edge (m,n) at time t . For link prediction, $f(m,n,t) \in \mathbb{R}^2$ is also a one-hot vector, but now encoding if the edge is an existing or non-existing edge. As appropriate, we weigh the minority class more heavily in the loss function for skewed datasets, and treat α as a hyperparameter. See Section B of the supplement for further details on the experiment setup, including the training setup and how hyperparameter tuning is done.

The experiments are implemented in PyTorch with some preprocessing done in Matlab. Our code is available at <https://github.com/IBM/TM-GCN>. In the experiments, we use an edge life of $l = 10$, a bandwidth $b = 20$, and $F' = 6$ output features. For TM-GCN, we

consider two variants of the \mathbf{M} matrix (M1 and M2); see the supplement for details.

We compare our method with three other methods. The first one is a variant of the WD-GCN by [19], which they specify in Equation (8a) of their paper. For the LSTM layer in their description, we use 6 output features instead of N . This is to avoid overfitting and make the method more comparable to ours which uses 6 output features. The second method is a 1-layer variant of EvolveGCN-H by [23]. The third method is a simple baseline which uses a 1-layer version of the GCN by [12]. It uses the same weight matrix \mathbf{W} for all temporal graphs. Both EvolveGCN-H and the baseline GCN use 6 output features as well. We use the prediction model (4.2) as the final layer in all models we compare.

Tables 2 and 3 show the results for edge classification and link prediction, respectively. For edge classification, our method outperforms the other methods on the two bitcoin datasets and the chess dataset, with WD-GCN performing best on the Reddit dataset. For link prediction, our method outperforms the other methods on the SBM, bitcoin and chess datasets. For Reddit, our method performs worse than the other methods. Results from some additional experiments are provided in Section C of the supplement.

5.3 COVID-19 Application One of the primary challenges related to the COVID-19 pandemic has been the issue of identifying early the individuals who are infected (ideally before they display symptoms) and prescribe testing. Here, we demonstrate how we can potentially use GNNs on contact tracing data to achieve this.

Contact tracing, a process where interactions between individuals (infected and others) are carefully tracked, has been shown to be an effective method for managing the spread of COVID-19. A variety of contact tracing methodologies have been used today around the world, see [1, 28] for lists. Recently, Ubaru et. al [28] presented a probabilistic graphical SEIR epidemiological model (Susceptible, Exposed, Infected, Recovered) to describe the dynamics of the disease transmission. Their model considers a dynamic graph (with individuals as nodes) that accounts for the interactions between individuals obtained from contact tracing, and uses a stochastic diffusion-reaction model to describe the disease transmission over the graph.

The novel SEIR model in [28] considers the graph Laplacian \mathbf{L}_t (from contact tracing data) at each time t and describes the evolution of the state $\{S, E, I, R\}$ for each node/individual. Figure 3 illustrates the state $\{S, E, I, R\}$ evolution as defined by the model on a sample dynamic graph with 10 individuals (for easy




Figure 3: Graphical SEIR model disease transmission visualization.

Table 4: COVID-19 Data: Mean absolute error and error ratio for infection state I prediction.

Methods	COVID-19 Dataset	
	Error	Ratio
WD-GCN	1.667	0.337
EvolveGCN	4.969	0.912
TM-GCN	1.466	0.278

visualization). We see how the infection (one individual as red node in first graph) transmits, we have magenta nodes with $I > 0.04$, and the yellow nodes with $I > 0.002$, and we note the interactions and the state change over time. Here, we show how we can use dynamic GNNs to predict the infection state I at time $T + 1$, using only the dynamic graphs up to time T , when the true SEIR model is unknown.

We consider a simulation with $N = 1000$ individuals and total time $T = 100$. We simulate the contact tracing dynamic graph as in [28], and assume at each time t a small number of individuals are tested at random for both IgM (if positive state I is set to 1) and IgG (state R is set to 1) antigen tests. The state of the remaining individuals are determined by the SEIR model. We train the dynamic GNNs on the first $T = 80$ time instances and test the GNNs on the remaining 20 time instances. Our goal is to train a GNN that learns the relation between the contact tracing graphs and the infection state (some exact values for those who were tested and others from the SEIR model), in order to *better predict* the individuals' state I at time $t + 1$ than just using the SEIR model. Table 4 gives the mean absolute error and error ratio obtained by the three dynamic GNNs for infection state I prediction on the test time instances. We note that, the proposed TM-GCN yields best results among the three methods, since it has a better time awareness (explicitly considers b previous time instances via the M-product) than others. Using such predictions, we can issue early warnings to individuals who are infected and prescribe testing.

6 Conclusion

We have presented a novel approach for dynamic graph embedding which leverages the tensor M-product framework. We used it for edge classification and link prediction in experiments on five datasets, where it performed competitively compared to state-of-the-art methods. We also demonstrated the method's effectiveness in an important application related to the COVID-19 pandemic. Future research directions include further developing the theoretical guarantees for the method, investigating optimal structure and learning of the transform matrix \mathbf{M} , using the method for other prediction tasks, and investigating how to utilize deeper architectures for dynamic graph learning.

Acknowledgments

We thank Stephen Becker and Lingfei Wu for helpful discussions and feedback. We also thank the reviewers for their insightful comments and suggestions. Kilmer was partially supported by a grant from IBM T.J. Watson and by Tufts T-Tripods Institute under NSF HDR grant CCF-1934553. Avron was supported by Israel Science Foundation grant 1272/17 and by an IBM Faculty Award.

References

- [1] H. Alsdurf, Y. Bengio, T. Deleu, P. Gupta, D. Ippolito, R. Janda, M. Jarvie, T. Kolody, et al. Covi white paper. *arXiv preprint arXiv:2005.08502*, 2020.
- [2] T. Y. Berger-Wolf and J. Saia. A framework for analysis of dynamic social networks. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 523–528. ACM, 2006.
- [3] K. Braman. Third-order tensors as linear operators on a space of matrices. *Linear Algebra and its Applications*, 433(7):1241–1253, 2010.
- [4] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun. Spectral networks and locally connected networks on graphs. *arXiv preprint arXiv:1312.6203*, 2013.

- [5] H. Cai, V. W. Zheng, and K. C.-C. Chang. A comprehensive survey of graph embedding: Problems, techniques, and applications. *IEEE Transactions on Knowledge and Data Engineering*, 30(9):1616–1637, 2018.
- [6] F. De Vico Fallani, J. Richiardi, M. Chavez, and S. Achard. Graph analysis of functional brain networks: Practical issues in translational neuroscience. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1653):20130521, 2014.
- [7] M. Defferrard, X. Bresson, and P. Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*, pages 3844–3852, 2016.
- [8] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1263–1272. JMLR. org, 2017.
- [9] P. Goyal, N. Kamra, X. He, and Y. Liu. Dyngem: Deep embedding method for dynamic graphs. *arXiv preprint arXiv:1805.11273*, 2018.
- [10] E. Kernfeld, M. Kilmer, and S. Aeron. Tensor-tensor products with invertible linear transforms. *Linear Algebra and its Applications*, 485:545–570, 2015.
- [11] M. E. Kilmer, K. Braman, N. Hao, and R. C. Hoover. Third-order tensors as operators on matrices: A theoretical and computational framework with applications in imaging. *SIAM Journal on Matrix Analysis and Applications*, 34(1):148–172, 2013.
- [12] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- [13] T. G. Kolda and B. W. Bader. Tensor Decompositions and Applications. *SIAM Review*, 51(3):455–500, 2009.
- [14] S. Kumar, W. L. Hamilton, J. Leskovec, and D. Jurafsky. Community interaction and conflict on the web. In *Proceedings of the 2018 World Wide Web Conference*, pages 933–943. International World Wide Web Conferences Steering Committee, 2018.
- [15] J. Kunegis. Konect: The koblenz network collection. In *Proceedings of the 22nd International Conference on World Wide Web*, pages 1343–1350. ACM, 2013.
- [16] J. Li, H. Dani, X. Hu, J. Tang, Y. Chang, and H. Liu. Attributed network embedding for learning in a dynamic environment. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pages 387–396. ACM, 2017.
- [17] Y. Li, R. Yu, C. Shahabi, and Y. Liu. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. *arXiv:1707.01926*, 2017.
- [18] X. Liu, X. You, X. Zhang, J. Wu, and P. Lv. Tensor graph convolutional networks for text classification. *arXiv preprint arXiv:2001.05313*, 2020.
- [19] F. Manessi, A. Rozza, and M. Manzo. Dynamic graph convolutional networks. *Pattern Recognition*, 97, 2020.
- [20] E. Mavroudi, B. B. Haro, and R. Vidal. Neural message passing on hybrid spatio-temporal visual and symbolic graphs for video understanding. *arXiv preprint arXiv:1905.07385*, 2019.
- [21] E. Newman, L. Horesh, H. Avron, and M. Kilmer. Stable Tensor Neural Networks for Rapid Deep Learning. *arXiv preprint arXiv:1811.06569*, 2018.
- [22] G. H. Nguyen, J. B. Lee, R. A. Rossi, N. K. Ahmed, E. Koh, and S. Kim. Continuous-time dynamic network embeddings. In *Companion Proceedings of the The Web Conference 2018*, pages 969–976, 2018.
- [23] A. Pareja, G. Domeniconi, J. Chen, T. Ma, T. Suzumura, H. Kanezashi, T. Kaler, T. B. Schardl, and C. E. Leiserson. EvolveGCN: Evolving graph convolutional networks for dynamic graphs. *arXiv preprint arXiv:1902.10191*, 2019.
- [24] A. Sankar, Y. Wu, L. Gou, W. Zhang, and H. Yang. Dynamic graph representation learning via self-attention networks. *arXiv preprint arXiv:1812.09430*, 2018.
- [25] Y. Seo, M. Defferrard, P. Vandergheynst, and X. Bresson. Structured sequence modeling with graph convolutional recurrent networks. In *Conference on Neural Information Processing*, pages 362–373. Springer, 2018.
- [26] R. Trivedi, H. Dai, Y. Wang, and L. Song. Know-evolve: Deep temporal reasoning for dynamic knowledge graphs. In *International Conference on Machine Learning-Volume 70*, pages 3462–3471. JMLR. org, 2017.
- [27] R. Trivedi, M. Farajtabar, P. Biswal, and H. Zha. DyRep: Learning representations over dynamic graphs. In *International Conference on Learning Representations*, 2019.
- [28] S. Ubaru, L. Horesh, and G. Cohen. Dynamic graph based epidemiological model for COVID-19 contact tracing data analysis and optimal testing prescription. *arXiv preprint arXiv:2009.04971*, 2020.
- [29] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon. Dynamic graph CNN for learning on point clouds. *arXiv preprint arXiv:1801.07829*, 2018.
- [30] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu. A comprehensive survey on graph neural networks. *arXiv preprint arXiv:1901.00596*, 2019.
- [31] T. Zhang, W. Zheng, Z. Cui, and Y. Li. Tensor graph convolutional neural network. *arXiv preprint arXiv:1803.10071*, 2018.
- [32] L. Zhao, Y. Song, C. Zhang, Y. Liu, P. Wang, T. Lin, M. Deng, and H. Li. T-GCN: A Temporal Graph Convolutional Network for Traffic Prediction. *IEEE Transactions on Intelligent Transportation Systems*, 2019.
- [33] J. Zhou, G. Cui, Z. Zhang, C. Yang, Z. Liu, and M. Sun. Graph neural networks: A review of methods and applications. *arXiv preprint arXiv:1812.08434*, 2018.
- [34] L. Zhou, Y. Yang, X. Ren, F. Wu, and Y. Zhuang. Dynamic network embedding by modeling triadic closure process. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.