

## Graph Neural Networks for Traffic Forecasting

João Rico\*<sup>1,2,3</sup>, José Barateiro<sup>1,2</sup>, Arlindo Oliveira<sup>2,3</sup>



<sup>1</sup>Laboratório Nacional de Engenharia Civil, Av. do Brasil 101, Lisbon, Portugal

<sup>2</sup>INESC-ID, R. Alves Redol 9, Lisbon, Portugal

<sup>3</sup>Instituto Superior Técnico, University of Lisbon, Lisbon, Portugal jmrigo@lnec.pt,

jbarateiro@lnec.pt, arlindo.oliveira@tecnico.ulisboa.pt

### Abstract:

The significant increase in world population and urbanisation has brought several important challenges, in particular regarding the sustainability, maintenance and planning of urban mobility. At the same time, the exponential increase of computing capability and of available sensor and location data have offered the potential for innovative solutions to these challenges. In this work, we focus on the challenge of traffic forecasting and review the recent development and application of graph neural networks (GNN) to this problem. GNNs are a class of deep learning methods that directly process the input as graph data. This leverages more directly the spatial dependencies of traffic data and makes use of the advantages of deep learning producing state-of-the-art results. We introduce and review the emerging topic of GNNs, including their most common variants, with a focus on its application to traffic forecasting. We address the different ways of modelling traffic forecasting as a (temporal) graph, the different approaches developed so far to combine the graph and temporal learning components, as well as current limitations and research opportunities.

Keywords: graph neural networks, traffic forecasting, deep learning

### Introduction

The world's urban population will increase from 4.2 billion to 6.7 billion by 2050 as estimated by the United Nations [1]. In spite of the accompanying social evolution and benefits, the rapid rate of urbanization has significant social, economic and environmental costs associated, including air and water pollution, unsustainable energy consumption, toxic waste disposal, inadequate urban planning, decreased public health and safety, social vulnerability and community disruption. Notably, the mobility of passengers and freights in most large cities of the world is not yet sustainable. Urbanization has given rise to traffic congestion, increase of transports needs, ineffective accessibility, and reduced productivity. In particular, traffic congestion costs billions of dollars per year due to lost time, air pollution, and wasted fuel. In 2017, in the United States alone, traffic congestion induced a total of

8.8 billion hours of travel delay and 12.5 billion liters of extra fuel consumption, corresponding to a congestion cost of 166 billion dollars [2]. Efforts to develop solutions to the challenges brought by traffic congestion have focused on three avenues: championing transport alternatives, enlarging the infrastructure, and managing traffic flows [3]. While championing transport alternatives is mostly a public policy issue, and geographical and social constraints limit the increase of the size of infrastructure, the potential to efficiently manage traffic flows has increasingly become one important solution to traffic congestion. Today, the exponential increase of available data in cities and the growth of computing capabilities represents a crucial opportunity to tackle these challenges by leveraging innovative and



integrated solutions. These include the development of intelligent transportation systems (ITS), smart vehicle sharing systems and home automation, and smart grids and energy solutions - all of which fall under the umbrella of the recent area of urban computing [4].

A core component of ITS is traffic forecasting. Its goal is to measure, model and predict traffic conditions in real-time, accurately and reliably, in order to optimize the flow and mitigate the congestion of traffic, and to respond adequately to other problems such as traffic light control, time of arrival estimates, and planning of new road segments. However, this is a very challenging problem due to several important factors. The successful forecasting of traffic conditions requires adequately handling heterogeneous data (e.g., integrating loop counter and floating car data) with complex spatio-temporal dependencies which are typically sparse, incomplete and high-dimensional. In addition, it requires computing in real-time and the inclusion of external factors such as weather conditions and road accidents.

This review is focused on the recent developments and applications of graph neural networks, a new family of deep learning models, to road traffic forecasting. In Section 2, we present a succinct history of traffic forecasting including traditional approaches and deep learning models that have had a great deal of success. Section 3 presents and discusses graph neural networks (GNN) and Section 4 reviews the literature of traffic forecasting based on GNN. We conclude, in Section 5, with a discussion of open challenges and research opportunities.

## Traffic forecasting

With the weighted adjacency matrix. Denoting  $X_{i,t}$  as the values of all the features of node  $i$  at time  $t$ , we aim to learn a function  $F$  that given a sequence of  $M$  historical time steps predicts the next  $N$  time steps:

$$F([X_{1,t}, \dots, X_{i,t}]; G) = [X_{1,t+N}, \dots, X_{i,t+N}]. \quad (1)$$

## Traffic data

Urban data can be classified according to their structures and spatiotemporal properties [9], as Figure 1 illustrates. We can distinguish datasets with respect to their spatiotemporal properties: spatiotemporal static data, spatial static but temporal dynamic data, or spatiotemporal dynamic data. They can also be classified with respect to their data structures in one of two types - point-based and network-based. In addition, urban data can also be classified according to its sources, such as geographical data, traffic data, and social network data, to name a few, and each of these can be further subclassified. Focusing on traffic data, this source can include loop detector data, floating car data, mobile phone data, call detail records, surveillance cameras, bike-sharing data, taxi records, mobile phone location data, parking records, and mobile apps' logs.



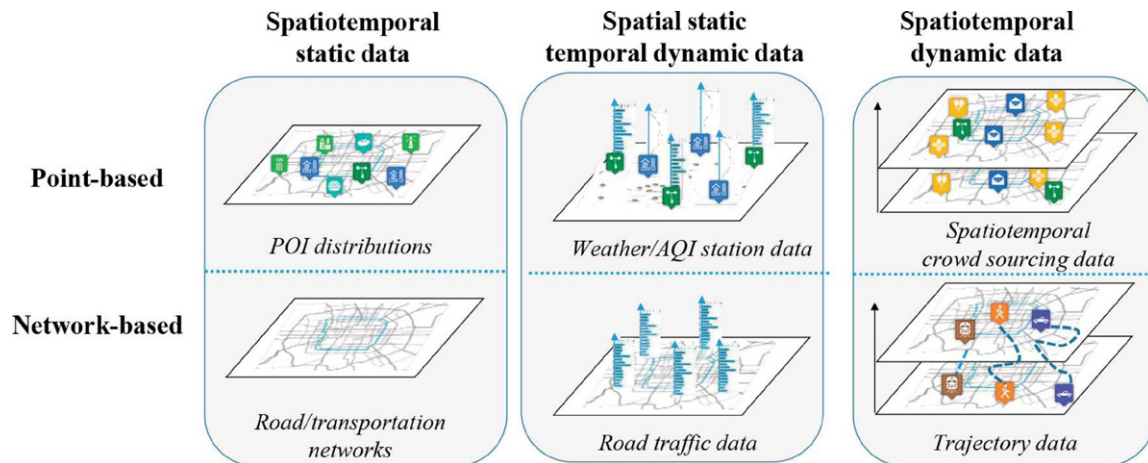


Figure 1: Six types of urban data [9]. Reprinted with permission.

### Traditional approaches to traffic forecasting

In the last decades, traffic forecasting methods have been developed by researchers from different fields such as transportation systems [10], economics [11], statistics, and machine learning [12]. The methods can be divided into two categories: the knowledge- or model-driven approach and the data-driven approach. Knowledge-driven approaches usually aim at modeling and explaining the transportation network through differential equations and numerical simulation [13], [14]. Although the models can reproduce real traffic conditions fairly accurately, they require prior knowledge and detailed modeling, are not easily transferable to other cases, and require significant computational resources.

Traditional data-driven methods used in traffic forecasting can be divided into two categories: methods that do not model the spatial dependency and models that do model this dependency. The former category includes the Historical Average (HA) model, the auto-regressive integrated moving average (ARIMA) [15] and seasonal ARIMA (SARIMA) [16], K-nearest Neighbour (KNN) [17], support vector regression [18], hidden Markov models [19], among others. These approaches often require careful feature engineering and also rely on data to satisfy certain assumptions, such as stationarity. However, real data is often too complex and violates often these assumptions, leading in many cases to poor performance. Approaches that model the spatial dependency include Vector ARIMA, Spatiotemporal ARIMA and Spatiotemporal HMM [20]–[23]. Again, these methods perform poorly since they are not complex enough to model the non-linearity and non-stationarity of the data.

### Deep learning approaches to traffic forecasting

Deep learning methods [24] are a class of machine learning methods that learn multiple layers of representations by composing increasingly more complex non-linear features on the upper layers by combining simpler features from the lower layers. This learning of complex representations is done mostly automatically and does not depend on a human doing manual feature engineering, which would require time and expert domain knowledge. In recent years, deep learning represents the state of the art in fields such as image recognition [25], natural language understanding [26],



drug discovery [27], recommendation systems [28], and board and video games playing [29]. This success is mainly due to their representation power (as explained above) and to the fact that there exists an efficient method for training them (namely gradient descent through error backpropagation).

As in other domains, the application of deep learning methods to traffic forecasting has been very successful and has produced state-of-the-art results. Some of the earlier architectures did not model the spatial dependency, and used standard feedforward networks or deep belief networks [30] as well as recurrent neural networks such as long short-term memory (LSTM) and the gated recurrent unit (GRU) [31]–[33]. However, these models still fail to model the complex spatial dependencies that exist in traffic problems. Several models were proposed that aim at capturing these dependencies, using convolutional neural networks (CNN) [34] or a combination of CNN and LSTM [35], [36]. Still, since CNNs are mainly suited for data embedded in grid-line Euclidean spaces, they are not the natural architecture to real-world road networks. The next sections expand on how models using a graph neural network improved on these architectures.

## Graph neural networks

The convolutional operator used in CNNs is very powerful but is limited to standard grid data, i.e., to data that originates in regular, two-dimensional, tri-dimensional, or higher-dimensional Euclidean spaces. Since various important machine learning problems involve tasks on graph structured data such as node classification [37] or molecule generation [38], researchers have developed a family of deep learning models that can leverage this inductive bias [37], [39]–[41]. These models are usually called graph neural networks (GNN) because they are neural networks that naturally handle graph data. In this section, we briefly review the area of graph neural networks and some of the most relevant works, and refer the reader to other reviews for more details and specific perspectives on this topic [41]– [46]. In Section 3.1, we categorize and describe several GNN models, and in Section 3.2 we list current open-source libraries for implementing GNNs.

## Models

In this section, we describe several of the most relevant GNNs models developed so far. As shown in Table 1, we categorize GNNs into four main types, namely recurrent graph neural networks, convolutional graph neural networks, graph autoencoders and graph attention networks.

Category	References
Recurrent Graph Neural Networks	[50]–[47] ,[39]
Convolutional Graph Neural Networks	[53]–[51] ,[40] ,[37]
Graph Attention Networks	[55] ,[54]
Graph Autoencoders	[61]–[56]

Table 1. Categorization of graph neural network models and representative publications.



## Recurrent graph neural networks

Extending previous work [62], [63], the Graph Neural Network model (GraphNN) [39] was the first neural network model that could process general types of graphs (eg, directed, undirected, cyclic, or acyclic). The fundamental concept of the GraphNN model is that every node  $v$  can be represented by a low-dimensional state vector  $h_v$  and be defined by its features and by its neighbours based on an information diffusion mechanism from every node to its neighbors. The goal is to learn this representation which can be fed to an output function  $g$ , called the local output function, resulting in an output value or label  $o_v$ , for regression or classification, respectively. The model also defines the local transition function  $f$ , a parametric function, to be learned alongside  $g$  - both shared by all nodes. Together, these node representation and output are defined as:

$$h_v = f(x_v, x_v^{[e]}, h_v^{[e]}, x_v^{[e]}), \quad (1)$$

$$o_v = g(h_v, x_v), \quad (2)$$

where  $x_v$ ,  $x_v^{[e]}$ ,  $h_v^{[e]}$ ,  $x_v^{[e]}$  are the features of  $v$ , the features of its edges, and the states and features of its neighborhood, respectively.

Denoting the vectors constructed by stacking all the states, all the features, all the node features and all the outputs by

$H, X, X^e, O$  equations (1) and (2) can be rewritten compactly as

$$H = F(H, X), \quad (3)$$

$$O = G(H, X^e), \quad (4)$$

where  $F$  and  $G$  are the stacked versions of  $f$  and  $g$  for all nodes, respectively.

If  $F$  is a contraction map, Banach's Fixed Point Theorem [64] guarantees the existence and uniqueness of the system of equations (3) and (4). This suggests iteratively updating the following equation:

$$H^t = F(H^{t-1}, X) \quad (5)$$

where  $H^t$  is the  $t$ -th iteration of  $H$ , and  $H^0$  can be initialized randomly.



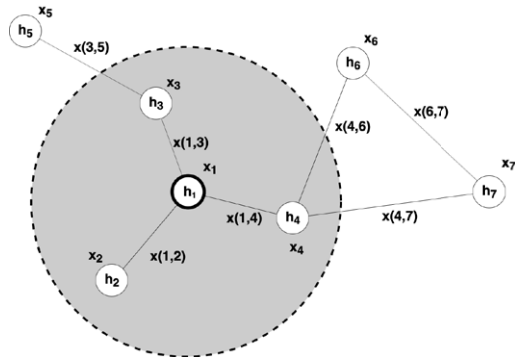


Figure 2 shows the intuition behind equations (1) and (5): to learn a node representation by propagating information from its neighbors, iteratively, until convergence.

$$h^i = (x^i, x^{(i,1)}, x^{(i,2)}, x^{(i,3)}, h^i, h^i, h^i, x^i, x^i, x^i)$$

Figure 2: Example graph of how GraphNN propagates information from the states and features of the immediate neighborhood of a node.

To learn the functions  $f$  and  $g$ , GraphNN defines a loss function at the supervised nodes, and uses a gradient descent

scheme: iteratively updating  $H^i$  until a convergence criterion is reached; calculating the gradients of the loss with respect to the weights of  $f$  and  $g$ ; and updating the weights.

Although it is a simple and powerful model, GraphNN has several limitations. Most importantly, it is computationally inefficient to compute the iterations for the fixed point and requiring  $F$  to be a contraction map limits the modeling capacity of the approach, including the long range dependencies of nodes (see appendix A of [47]).

The Gated Graph Neural Network (GatedGNN) [47] is another recurrent graph neural network, and it improves on some of the drawbacks of GraphNN. This model modifies the original GraphNN substituting the recurrence function in equation (1) for a Gated Recurrent Unit (GRU) [65]. GatedGNN uses backpropagation through time (BPTT), unrolling the recurrence for a fixed number of  $T$  steps. While for a large graph this can present a drawback by requiring a large amount of memory to store the intermediate states computed over all nodes, on the other hand this removes the requirement for a contraction map in order to guarantee convergence. The model recurrence is defined as follows



$$a^{!!} = A^{!!}[h^{!!!!} \dots h^{!!!!}] + b, \quad g^{!!} = \tanh(Wa^{!!} + U(h^{!!!!} \odot h^{!!!!})),$$

$$z^{!!} = \sigma(W'a^{!!} + U'h^{!!!!}), \quad h^{!!} = (1 - z^{!!}) \odot h^{!!!!} + (z^{!!} \odot g^{!!}), \quad (6)$$

$$r^{!!} = \sigma(W'a^{!!} + U'h^{!!!!}),$$

where  $A$  is the modified adjacency matrix composed of two submatrices, one for outgoing edges, the other for incoming edges. Figure 3 shows the parameter tying and sparsity structure in the adjacency matrix, as well as an example of the unrolling of the recurrence for one time step.  $Av$ : corresponds to the two columns of the submatrices  $A^{(out)}$  and  $A^{(in)}$  referring to node  $v$ , and  $z$  and  $r$  correspond to the update and reset gates of the GRU.

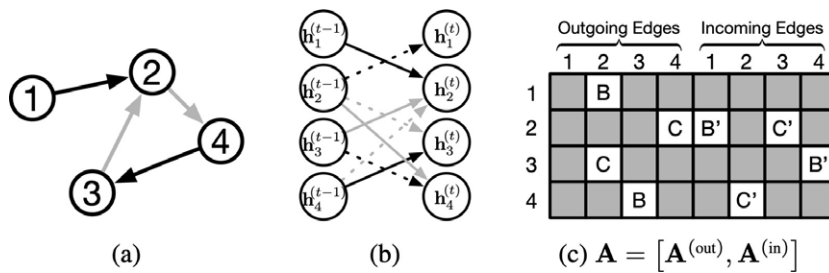


Figure 3: (a) Example graph, where edge color denotes edge type. (b) Unrolling of one time step. (c) Parameter tying and sparsity in adjacency matrix. Letters denote edge types.  $B'$  corresponds to the reverse edge of type  $B$ .  $B$  and  $B'$  denote distinct parameters [47]. Reprinted with permission.

Other recurrent graph neural networks have been developed, including the Graph Echo State Network [66], Graphrnn

[50] and Stochastic Steady-state Embedding [49].

## Convolutional graph neural networks

The generalization of the very successful CNN to non-Euclidean data such as graphs and manifolds has followed two distinct avenues, the so-called spectral methods and spatial methods. Briefly, spectral based methods define the convolution based on the graph Laplacian and use the corresponding Fourier basis in the spectral domain, and spatial methods use a localized parameter-sharing filter on the neighbors of each node, in essence, aggregating local features.

Spectral Convolutional Neural Networks (SpectralCNN) [51] was the first model to introduce the convolution for graph data. Given an undirected graph its normalized graph Laplacian is defined as



$$L = I - D^{-1}AD, \quad (7)$$

where  $A$  is the adjacency matrix and  $D$  is the diagonal matrix of node degrees. The graph Laplacian is a real symmetric positive semidefinite matrix which implies that it can be factorized as

$$L = \Lambda U^T, \quad (8)$$

where  $\Lambda$  is the diagonal matrix of eigenvalues - the spectrum - and  $U$  is the matrix of eigenvectors. Given a graph signal

$x \in \mathbb{R}^n$  where  $x_i$  is the feature of the  $i$ -th node. The graph Fourier transform  $\mathcal{F}$  - and its inverse  $\mathcal{F}^{-1}$  are defined by

$$\mathcal{F}(x) = U^T x = \lambda, \quad (9)$$

$$\mathcal{F}^{-1}(\lambda) = U \lambda. \quad (10)$$

The graph convolution operation of a graph signal  $x$  with a filter  $g$  on a graph  $G$  is defined as

$$x * g = \mathcal{F}^{-1}(\mathcal{F}(x) \odot \mathcal{F}(g)),$$

$$= U(U^T x \odot U^T g) \quad (11)$$

where  $\odot$  is the Hadamard product. The equation above can be simplified, by considering a filter  $g = \text{diag}(\theta)$  parametrized by  $\theta \in \mathbb{R}^n$  in the Fourier domain, to the following expression

$$x * g = U g U^T x, \quad (12)$$

where filter  $g$  can be understood as a function of eigenvalues of the Laplacian, that is  $g(\Lambda)$ .

This model has several drawbacks. It is computationally inefficient since multiplication by  $U$  is  $O(n^2)$  and calculating the eigen-decomposition is  $O(n^3)$ . In addition, the dimensionality of the trainable filter depends on the number of nodes of the graph, and the filters are non-localized depending on the graph structure - if a graph changes at all, so will its eigenvalues and eigenvectors.

Chebyshev Spectral Convolutional Neural Networks (ChebNet) [40] proposes to approximate the filter by Chebyshev polynomials of the diagonal matrix of eigenvalues, that is,

$$h_i = \sigma \left( \sum_{l \in \{0, 1, \dots, L\}} \alpha_l W^{(l)} h_i^{(l)} \right),$$





where  $L^* = 2L/\lambda^* - I$  and  $\lambda^*$  is the maximum eigenvalue. The Chebyshev polynomials are defined by the recursion  $T_k(x) = 2xT_{k-1}(x) - T_{k-2}(x)$  with  $T_0(x) = 1$  and  $T_1(x) = x$ . In this setting, the computational complexity reduces to  $O(KM)$  where  $K$  is the polynomial order of the expansion above and  $M$  is the number of edges, rendering the model linear with respect to the graph size. This also renders the filters spatially localized, since for an expansion of degree  $K$  each node only receives information from a node at maximum  $K$  hops away.

Graph Convolutional Networks (GCN) [37] extends ChebNet by limiting the Chebyshev expansion to only the first order,  $K = 1$ , and employing some additional modifications. It assumes  $\lambda^* = 2$ , simplifying equation (13) to

$$x * g = \theta x - \theta D^{-1} A D^{-1} x. \quad (14)$$

Setting  $\theta = -\theta$  and substituting  $A$  for  $A^* = A + I$ , equation (12) can be written as

$$x * g = \theta (I + D^{-1} A^* D^{-1}) x. \quad (15)$$

Generalizing to the signal to  $C$  input channels and  $F$  filters for feature maps, the graph convolution can be written as

$$Z = D^{-1} A^* D^{-1} \theta, \quad (16)$$

where  $Z \in R^{N \times C}$ , and  $\theta \in R^{C \times F}$  is a matrix of filter parameters.

ChebNet and GCN bridge the gap between spectral- and spatial based methods. Equation (16) can be expressed as

$$\sum_{(16)}$$

$$h_i = f(($$

$$i \in \{i | U_i$$

$$A_i(x) \theta),$$

which can be interpreted as a spatial-based method, each node aggregating information from its neighbors.

$\sum$  Neural Fingerprints (Neural FPs) [53] generalized GNNs by using different learning parameters for nodes with different degrees, as follows:

$$x = h^{(k)} +$$

$$!!!$$

$$h^{(k)} \theta T_k(A^*) U x,$$



(17)

$$h^{(t)} = \sigma(xW^{(t)}), \quad (18)$$

where  $W^{(t)}$  is the weight matrix for nodes of degree  $N$  at layer  $t$ .

Another interesting model is the dual graph convolutional network (DGCN) [67] which jointly uses two convolutions: one based on the adjacency matrix is the same as equation (16), and the second one is based on a diffusion process to capture global relationships by substituting the adjacency matrix  $A$  with the positive pointwise mutual information (PPMI) matrix [52].

## Graph Attention Networks

Attention mechanisms [68], [69] are a recent development in deep learning that has had success in tasks such as machine translation. Figure 4 illustrates the intuition of why attention mechanisms could be helpful in modelling the neighborhood of a node, namely by leveraging the type of each neighbor to assign different importance weights.

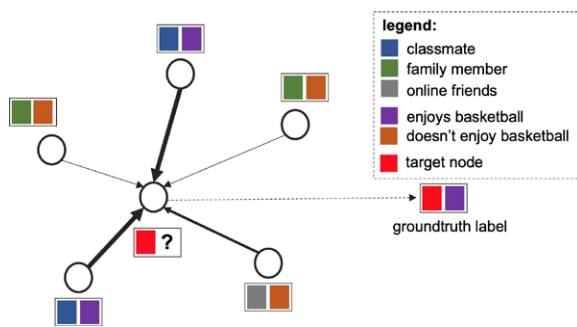


Figure 4: The type of each neighbor is used to assign attention, and the link size denotes how much attention to apply to each neighbor. This example illustrates how concentrating our attention on the node's classmates, would improve our prediction of the enjoyment of the activity [70]. Reprinted with permission.

Graph Attention Networks (GAT) [54] was the first model to incorporate an attention mechanism in GNNs by removing the requirement that all the neighbors of a node contribute with equal or pre-defined weights. Rather, the relative weights between two connected weights are learned as follows:

$$h^{(t)} = \sigma\left(\sum_{i \in \mathcal{N}(j)} \alpha_{ij} W^{(t)} h^{(t-1)}_i\right),$$



(19)

where  $\alpha^i$  is the attention coefficient of node  $j$  to  $i$  defined as

$$\alpha^i = \text{softmax}(g(\mathbf{a}^T [W^i h^j || W^i h^i])), \quad (20) \text{ where } \mathbf{a} \text{ is a vector of learnable parameters and } g \text{ is a LeakyReLU unit.}$$

GAT also proposes incorporating a multi-head attention mechanism [69], and have the important advantage of being efficient since the computation of the node-neighbor pairs is parallelizable and because it can naturally be applied to nodes with a different number of neighbors.

Gated Attention Network (GAAN) [55] is an extension of GAT which, in addition to the multi-head attention mechanism, also incorporates a self-attention mechanism which replaces the averaging operation of GAT with an attention score on each attention head. For more details and different taxonomies of attention models in graphs, we refer the reader to a comprehensive review [70].

## Graph Autoencoders

Graph neural networks have also been used in unsupervised learning settings. In particular, successful deep learning models such as autoencoders [71] and variational autoencoders [72] have been generalized to handle graph data. These models are usually developed to either learn network embeddings, such as Structural Deep Network Embedding [56] and Variational Graph Autoencoder [57], or to generate new graphs such Deep Generative Model of Graphs [58] and Graph Variational Autoencoder [59].

Some models have also incorporated adversarial training, in particular Generative Adversarial Networks (GAN) [73]. In a GAN, two networks compete against each other: the generator network is optimized to produce realistic samples of data that the discriminator network tries to distinguish from real data. GANs have been used together with GCN for generation of molecular graphs [60], link prediction, node clustering, and graph visualization [61].

## Frameworks

In addition to the different types of GNN models that have been developed so far, some of which we describe above, there has been an effort to unify GNNs under a general and common framework. Non-local Neural Networks (NLNN)

[74] unifies different self-attention mechanisms, and Message Passing Neural Networks (MPNN) [75] generalizes several GNN approaches using a message passing scheme. In MPNNs, messages  $m^t$  at timestep  $t$  are passed from each node  $v$  as follows:

$$m^t =$$

$$\sum$$

$$u \in \mathcal{N}(v)$$

$$M^t(h^v, h^u, e^v),$$



(21)

$$h^{v,u} = U^v(h^v, m^{v,u}), \quad (22)$$

where  $U^v$  is the vertex update function,  $M^v$  is the message function and  $e^{v,u}$  represents the features of the edge from node

v to u.

The Graph Networks (GN) model [41] generalizes both the NLNN and MPNN frameworks and other types of GNNs. It defines a GN block containing three update functions and three aggregation functions as follows:

$$e^{v'} = \varphi^v(e^v, h^v, h^u, u), \quad e^{*v'} = \rho^v \rightarrow^v(E^v),$$

$$h^{v'} = \varphi^v(e^{*v'}, h^v, u), \quad e^{*v'} = \rho^v \rightarrow^v(E^v), \quad (23)$$

$u^v = \varphi^v(e^{*v'}, h^{*v'}, u)$ ,  $h^{*v'} = \rho^v \rightarrow^v(H^v)$ , where  $\rho^v()$  are the message passing functions and  $\varphi^v()$  the update functions.

Software

The research and application of GNNs has been greatly enhanced by the development and publishing of several open- source libraries, specifically written to handle GNNs. These are built on top of already powerful deep learning frameworks, such as PyTorch [76] and Tensorflow [77]. These GNN libraries can greatly reduce the time of testing and deployment of new models, by providing useful abstractions that simplify the code required while automatically handling several low-level optimizations, such as scaling, parallelization and taking advantage of sparse structure. It is also common for authors to release implementations of their models in these libraries, which are, in turn, often integrated in these libraries in a ready-to-use manner.

Library	Refer- ence	DL Framework	URL
DGL	[78]	MXNet, Pytorch	<a href="https://github.com/dmlc/dgl">https://github.com/dmlc/dgl</a>
Euler	[79]	Tensorflow	<a href="https://github.com/alibaba/euler">https://github.com/alibaba/euler</a>
Graph Nets	[41]	Tensorflow	<a href="https://github.com/deepmind/graph_nets">https://github.com/deepmind/graph_nets</a>
PyTorch Geometric	[80]	Pytorch	<a href="https://github.com/rusty1s/pytorch_geometric">https://github.com/rusty1s/pytorch_geometric</a>

Table 2: Graph deep learning libraries.



## Graph neural networks for traffic forecasting

The connection between graphs and roads or maps is as old as graph theory, dating back to Euler's formulation and solution to the Königsberg problem. Nodes can represent road intersections and edges can model the road segments between them. Or, almost as naturally, nodes can represent points on a road segment where a traffic variable is being measured by a sensor and edges represent some relationship between these locations, such as the shortest path distance between them. The addition of a temporal component can be just as natural, as Figure 6 illustrates. Typically, each node in a graph would define a time-series of a feature vector, such as the history of traffic speed and traffic flow at a certain location.

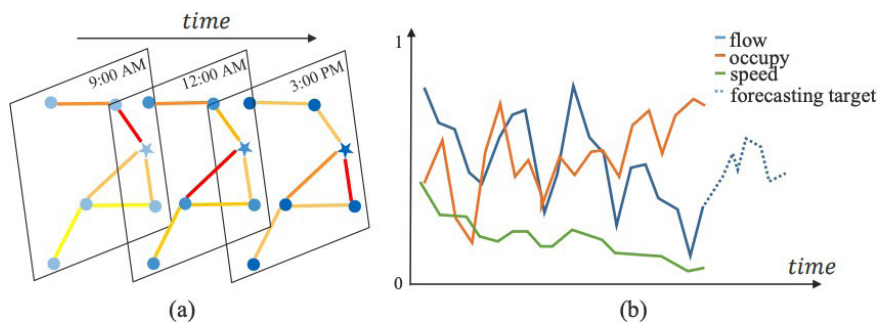


Figure 6: (a) Typical spatio-temporal structure of traffic data with each time slice forming a graph. (b) Representation of a three-dimensional time series at a certain node and a forecast of one of these features. [8]. Reprinted with permission.

Traffic forecasting based on graph neural networks can be viewed as the problem of extending the models of the previous section and the categories of Table 1 to the temporal domain. In principle, combining the many types of GNNs approaches with the immense diversity of existing time series analysis models and forecasting methods offers a wide space of possibilities.

In this section, we will present and compare the graph neural network models for traffic forecasting that have been developed so far. These have in common a graph neural network as a centerpiece of the model and they vary along most modern deep learning paradigms, combining recent techniques that have been found powerful in other domains of applications. These include attention mechanisms [8], [55], multi-graphs [81], dynamic graphs [82], diffusion kernels [83], graph wavenet [84], inception models [85], deep graph infomax [86], residual nets [87], wavelet transforms [88], among others [6], [7], [89]–[95]. In section 4.1. we review these models and compare different approaches, and in Section 2 we describe typical datasets in which these models are evaluated.



## Models

Table 3 lists the GNNs models developed for traffic forecasting so far, which we describe below. We start by describing Spatio-Temporal Graph Convolutional Networks (STGCN) [92] and Diffusion Convolutional Recurrent Neural Networks (DCRNN) [83]. As of the writing of this review, these two models stand out as the most cited and most often compared to when developing a new model, in part because they are some of the models first developed while also representing opposite approaches in some aspects.

STGCNs combine the spectral-based ChebNets [40] with 1D-CNNs to model and predict traffic speed at various locations where sensors are located. The GCN and the CNN operate alternately along several layers, with the GCN capturing the spatial dependency while the CNN captures the temporal dependency. Figure 7 shows an illustration of scheme similar to the one used by STGCN, in which a CNN is combined with a GCN.

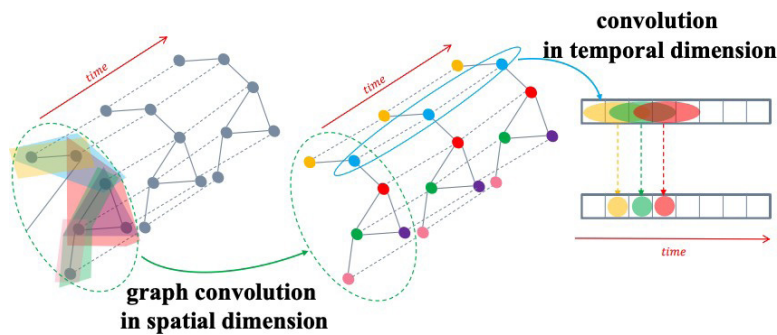


Figure 7: Illustration of how to combine GNN and CNN to capture the spatial and temporal dependencies, respectively [8]. Reprinted with permission.

DCRNNs take the opposite direction: they use RNNs instead of 1D-CNN to model the time dependencies and, in addition, use a spatial-based GCN for the spatial dependency instead of a spectral-based one. This is done by incorporating the GCN inside a GRU. Based on Diffusion-Convolution Neural Networks [96], DCRNN uses the probability transition matrix  $P = D^{-1}A$  to define the graph convolution as follows:

$$H =$$

!

$$\sum$$

!!!

$$f(P!XW(!)),$$



(24)

where  $f()$  is an activation function and  $W(!) \in R^{! \times !}$ . Effectively, DCRNN assumes that by passing information from node to node the network can reach an equilibrium state which is captured by the probability transition matrix. DCRNN also uses an encoder-decoder architecture to predict future timesteps.

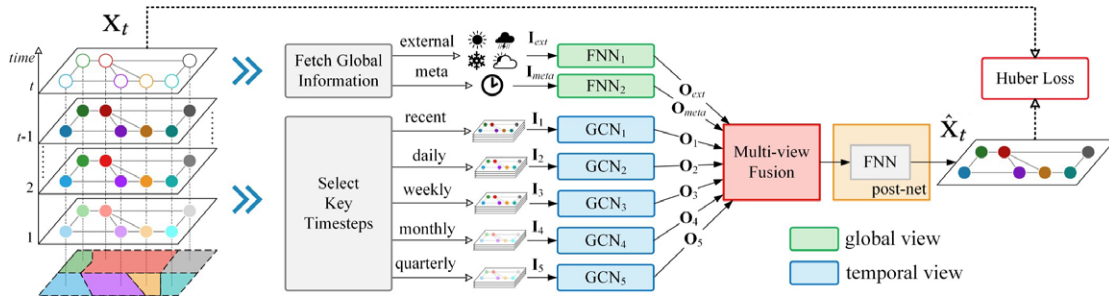


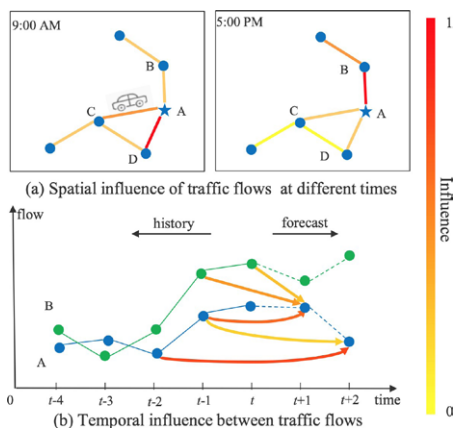
Figure 8: Architecture of MVGCN, displaying the views capturing different timescale dependencies, and their fusion in the final layer. [90]

A different and promising approach has been proposed by several recent models, including Multi-View Graph Convolutional Networks (MVGCN) [90], Multi Residual Recurrent Graph Neural Networks (MRes-RGNN) [87], Motif-based Graph Convolutional Recurrent Neural Network (Motif-GCRNN) [88] and Spatial-Temporal Graph Inception Residual Networks (STGI-ResNet) [85]. In these approaches, typical repeating time patterns are grouped in order to leverage the inductive bias we know is present in typical traffic, such as daily or weekly periods. As an illustration, if one is predicting traffic speed at 9am on a friday, in principle, one can make a better prediction by considering not only the last previous hours, but also, for example, the traffic at 9am in the previous four days of the week, and in the previous fridays, at 9am. These different views are then fused in the last layers of the architecture. This fusion layer is also learnable. Figure 8 shows the architecture of MVGCN which exemplifies this approach. As shown in the figure, these approaches can also typically handle exogenous factors such as weather and irregular events (e.g., road accidents), that have a significant impact in urban traffic.



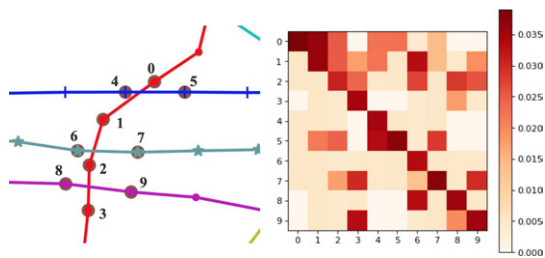
Model	.Ref	Scope	Predicts	Data source	Datasets	?Open dataset	Code avail- ?able
ST-GCN	[92]	Fw, Ur	S	L	BJER4, PeMS	✓, X	✓
DCRNN	[83]	Fw	S	L	METR-LA, PeMS	✓	✓
MRes-RGNN	[87]	Fw	S	L	METR-LA, PeMS	✓	X
TGC-LSTM	[7]	Fw, Ur	S	L, FCD	LOOP, INRIX	✓, X	X
ASTGCN	[8]	Fw	F, S	L	PeMSD4, PeMSD8	✓	✓
STDGI	[86]	Fw	S	L	METR-LA	✓	✓
MVGCN	[90]	Ur	F	FCD	TaxiNYC, Tax- iBJ, BikeDC, BikeNYC	✓	X
DST-GCNN	[82]	Fw, Ur	S, V	L, FCD	METR-LA, TaxiBJ	✓	X
GSRNN	[91]	Ur	F	FCD	BikeNYC, TaxiBJ	✓	X
Graph Wavenet	[84]	Fw	S	L	METR-LA, PeMS	✓	✓
3D-TGCN	[6]	Fw	S	L	PeMS	✓	X
ST-UNet	[93]	Fw	S	L	METR-LA, PeMS	✓	X
GaAN	[55]	Fw	S	L	METR-LA	✓	X
Mo- tif-GCRNN	[88]	Ur	S	FCD	TaxiChengdu	X	X
STGi-Res- Net	[85]	Ur	F	FCD	Didi Chengdu	✓	X
T-GCN	[94]	Fw, Ur	S	FCD	SZ-taxi, Los- loop	✓, X	X
FlowConv- GRU	[97]	Ur	F	FCD	TaxiNYC, TaxiCD	✓	X

Table 3: Graph Neural Networks approaches to traffic forecasting. For each publication, we list the scope of application (Ur: urban, Fw: freeway), the variables predicted (S: speed, F: flow, V: volume), the data source (L: loop counters, FCD: floating car data), the datasets used for experiments, whether or not these datasets are open, and whether there exists open-source code implementations of the models.





Attention mechanisms have also shown to be a promising approach to spatio-temporal forecasting. In addition to the GaAN architecture [55] described in Section 3.1.3, Attention Based Spatial-Temporal Graph Convolutional Networks (ASTGCN) [8] combine attention mechanisms on both the spatial and the temporal components with a multi-view approach on temporal patterns (namely, recent, daily-periodic and weekly-periodic). Figure 9 illustrates the intuition on how attention mechanisms can help make better predictions as well as output more interpretable models.



(c) Attention matrix obtained from the spatial attention mechanism.

Figure 9: (a) and (b) Spatial-temporal correlation diagram of traffic flow. (c) Attention matrix of sub-graph with 10 detectors, where the  $i$ -th row indicates how strong the time series of node  $i$  correlates with every other node. As an example of the interpretability aspect of this model, we can see that it indicates that pairs of nodes most strongly correlated tend to be spatially close, such as 1 and 6, 5 and 4, and 9 and 3. [8] Reprinted with permission.

Another rich approach has been to enable models to learn the graph adjacency matrix from the data and its patterns such as time series similarity between different nodes instead, or in addition to, the usual spatial distance. Variants of this approach are proposed by models that include Dynamic Spatio-Temporal Graph-based CNNs (DST-GCNN) [82], 3D Temporal Graph Convolutional Networks (3D-TGCN) [6], Graph WaveNet [84] and FlowConvGRU [97].

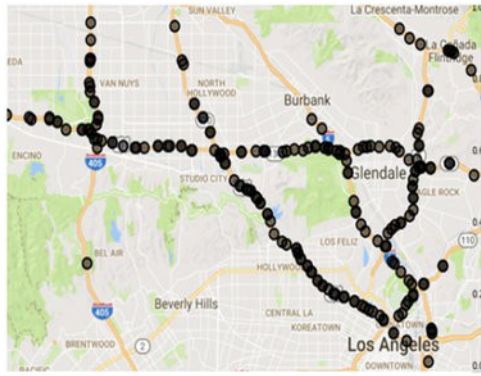
Finally, other models focus on leveraging other approaches such as the sparsity of traffic data [91], pooling and unpooling layers [93] or unsupervised learning of node representations [86].

## Datasets

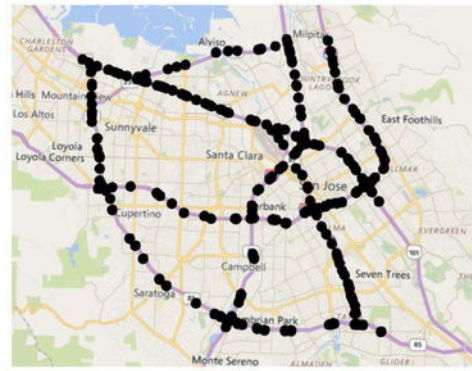
In spite of the differences described above, nearly every model has been experimentally evaluated in similar datasets with a similar configuration. In particular, the timestep of the datasets is, with few exceptions, 5 minutes, and the prediction horizon is 15, 30 and 60 minutes (that is, 3, 6, and 10 timesteps). The extension of the dataset is also very often around 4 months, which is equivalent to about 35000 timesteps.

Two datasets in particular stand out as the most used for benchmarking: METR-LA and PEMS-BAY. These datasets contain traffic information collected from loop sensors at various locations of two networks of freeways in California, one in Los Angeles, the other in the Bay Area. Typically only a small fraction of the data is used. For example, data is aggregated in 5 minutes interval and only between 200 and 1000 sensors are used to test the models.





(a) METR-LA



(b) PEMS-BAY

Figure 6: Sensor distribution of the METR-LA and PEMS-BAY datasets [83]. Reprinted with permission.

## Challenges and research opportunities

While the results obtained in various traffic forecasting tasks by deep learning models and, notably, graph neural networks have been very successful, there are still a number of open issues and traffic forecasting is still an unsolved problem. These challenges constitute a rich set of research and engineering opportunities [3], [98] which include integrating in a systematic way exogenous factors (such as road accidents and weather), designing more sophisticated evaluation metrics, integrating traffic forecasting with other downstream applications, going from volume or speed prediction to travel time prediction, improving the interpretability of the models and moving from prediction to causation. Additionally, ongoing work tackles the problem of developing Bayesian methods that provide adequate predictions beyond point estimates (such as confidence intervals) and tackling the issue of data ageing and concept drift by developing models that take into account when road segments or bus lines are added or removed from the network.

## Acknowledgments

This work was partly supported by the projects iLU-DSAIPA/DS/0111/2018 and UID/CEC/50021/2019 funded by the FCT (Fundação para a Ciência e Tecnologia). Work by João Rico is supported by a research grant by LNEC (Laboratório Nacional de Engenharia Civil) and FCT.



## References

- United Nations, Department of Economic and Social Affairs, Population Division., "World urbanization prospects. the 2018 revision," ST/ESA/SER. A/421, 2019.
- D. Schrank, B. Eisele, and T. Lomax, "2019 Urban Mobility Report," Texas A&M Transportation Institute, 2019.
- I. Lana, J. Del Ser, M. Velez, and E. I. Vlahogianni, "Road Traffic Forecasting: Recent Advances and New Challenges," *IEEE Intell. Transp. Syst. Mag.*, vol. 10, no. 2, pp. 93–109, 2018.
- Y. Zheng, L. Capra, O. Wolfson, and H. Yang, "Urban Computing: Concepts, Methodologies, and Applications," *ACM Trans. Intell. Syst. Technol.*, vol. 5, no. 3, pp. 38:1–38:55, Sep. 2014.
- E. I. Vlahogianni, M. G. Karlaftis, and J. C. Golias, "Short-term traffic forecasting: Where we are and where we're going," *Transp. Res. Part C: Emerg. Technol.*, vol. 43, pp. 3–19, 2014.
- B. Yu, M. Li, J. Zhang, and Z. Zhu, "3D Graph Convolutional Networks with Temporal Graphs: A Spatial Information Free Framework For Traffic Forecasting," *arXiv [cs.LG]*, 03-Mar-2019.
- Z. Cui, K. Henrickson, R. Ke, and Y. Wang, "Traffic Graph Convolutional Recurrent Neural Network: A Deep Learning Framework for Network-Scale Traffic Learning and Forecasting," *arXiv [cs.LG]*, 20-Feb-2018.
- S. Guo, Y. Lin, N. Feng, C. Song, and H. Wan, "Attention Based Spatial-Temporal Graph Convolutional Networks for Traffic Flow Forecasting," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019, vol. 33, pp. 922–929.
- Y. Zheng, *Urban Computing*. MIT Press, 2019.
- Smith Brian L. and Demetsky Michael J., "Traffic Flow Forecasting: Comparison of Modeling Approaches," *J. Transp. Eng.*, vol. 123, no. 4, pp. 261–266, Jul. 1997.
- K. A. Small, E. T. Verhoef, and R. Lindsey, *The economics of urban transportation*. Routledge, 2007.
- G. A. Davis and N. L. Nihan, "Nonparametric regression and short-term freeway traffic forecasting," *J. Transp. Eng.*, vol. 117, no. 2, pp. 178–188, 1991.
- D. Chowdhury, L. Santen, and A. Schadschneider, "Statistical physics of vehicular traffic and some related systems," *Phys. Rep.*, vol. 329, no. 4, pp. 199–329, May 2000.
- M. Saidallah, A. El Fergougui, and A. E. Elalaoui, "A Comparative Study of Urban Road Traffic Simulators," *MATEC Web of Conferences*, vol. 81, p. 05002, 2016.
- S. Lee and D. B. Fambro, "Application of Subset Autoregressive Integrated Moving Average Model for Short-Term Freeway Traffic Volume Forecasting," *Transp. Res. Rec.*, vol. 1678, no. 1, pp. 179–188, Jan. 1999.
- Williams Billy M. and Hoel Lester A., "Modeling and Forecasting Vehicular Traffic Flow as a Seasonal ARIMA Process: Theoretical Basis and Empirical Results," *J. Transp. Eng.*, vol. 129, no. 6, pp. 664–672, Nov. 2003.
- L. Zhang, Q. Liu, W. Yang, N. Wei, and D. Dong, "An Improved K-nearest Neighbor Model for Short-term Traffic Flow Prediction," *Procedia - Social and Behavioral Sciences*, vol. 96, pp. 653–662, Nov. 2013.
- H. Su, L. Zhang, and S. Yu, "Short-term Traffic Flow Prediction Based on Incremental Support Vector Regression," in *Third International Conference on Natural Computation (ICNC 2007)*, 2007, vol. 1, pp. 640–645.
- Y. Qi and S. Ishak, "A Hidden Markov Model for short term prediction of traffic conditions on freeways," *Transp. Res. Part C: Emerg. Technol.*, vol. 43, pp. 95–111, Jun. 2014.
- Y. Kamarianakis and P. Prastacos, "Forecasting Traffic Flow Conditions in an Urban Network: Comparison of Multivariate and Univariate Approaches," *Transp. Res. Rec.*, vol. 1857, no. 1, pp. 74–84, Jan. 2003.
- W. Min and L. Wynter, "Real-time road traffic prediction with spatio-temporal correlations," *Transp. Res. Part C: Emerg. Technol.*, vol. 19, no. 4, pp. 606–616, Aug. 2011.
- J. Kwon and K. Murphy, "Modeling freeway traffic with coupled HMMs," *Technical report*, Univ. California, Berkeley, 2000.
- J. Xu, D. Deng, U. Demiryurek, C. Shahabi, and M. van der Schaar, "Mining the Situation: Spatiotemporal Traffic Prediction With Big Data," *IEEE Journal of Selected Topics in Signal Processing*, vol. 9, no. 4, pp. 702–715, 2015.
- Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015.
- A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," in *Advances in Neural Information Processing Systems 25*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2012, pp. 1097–1105.
- R. Collobert, J. Weston, L. Bottou, M. Karlen, K. Kavukcuoglu, and P. Kuksa, "Natural Language Processing (Almost) from Scratch," *J. Mach. Learn. Res.*, vol. 12, no. Aug, pp. 2493–2537, 2011.
- J. Ma, R. P. Sheridan, A. Liaw, G. E. Dahl, and V. Svetnik, "Deep Neural Nets as a Method for Quantitative Structure-Activity Relationships," *Journal of Chemical Information and Modeling*, vol. 55, no. 2, pp. 263–274, 2015.
- H. Wang, N. Wang, and D.-Y. Yeung, "Collaborative Deep Learning for Recommender Systems," in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Sydney, NSW, Australia, 2015, pp. 1235–1244.
- V. Mnih et al., "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.



- W. Huang, G. Song, H. Hong, and K. Xie, "Deep Architecture for Traffic Flow Prediction: Deep Belief Networks With Multitask Learning," *IEEE Trans. Intell. Transp. Syst.*, vol. 15, no. 5, pp. 2191–2201, Oct. 2014.
- X. Ma, Z. Tao, Y. Wang, H. Yu, and Y. Wang, "Long short-term memory neural network for traffic speed prediction using remote microwave sensor data," *Transp. Res. Part C: Emerg. Technol.*, vol. 54, pp. 187–197, May 2015.
- R. Yu, Y. Li, C. Shahabi, U. Demiryurek, and Y. Liu, "Deep Learning: A Generic Approach for Extreme Condition Traffic Forecasting," in *Proceedings of the 2017 SIAM International Conference on Data Mining, Society for Industrial and Applied Mathematics*, 2017, pp. 777–785.
- Z. Cui, R. Ke, and Y. Wang, "Deep Bidirectional and Unidirectional LSTM Recurrent Neural Network for Network-wide Traffic Speed Prediction," *arXiv [cs.LG]*, 07-Jan-2018.
- X. Ma, Z. Dai, Z. He, J. Ma, Y. Wang, and Y. Wang, "Learning Traffic as Images: A Deep Convolutional Neural Network for Large-Scale Transportation Network Speed Prediction," *Sensors*, vol. 17, no. 4, Apr. 2017.
- J. Zhang, Y. Zheng, and D. Qi, "Deep spatio-temporal residual networks for citywide crowd flows prediction," in *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- H. Yu, Z. Wu, S. Wang, Y. Wang, and X. Ma, "Spatiotemporal Recurrent Convolutional Networks for Traffic Prediction in Transportation Networks," *Sensors*, vol. 17, no. 7, Jun. 2017.
- T. N. Kipf and M. Welling, "Semi-Supervised Classification with Graph Convolutional Networks," *arXiv [cs.LG]*, 09-Sep-2016.
- Q. Liu, M. Allamanis, M. Brockschmidt, and A. Gaunt, "Constrained Graph Variational Autoencoders for Molecule Design," in *Advances in Neural Information Processing Systems 31*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds. Curran Associates, Inc., 2018, pp. 7795–7804.
- F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini, "The graph neural network model," *IEEE Trans. Neural Netw.*, vol. 20, no. 1, pp. 61–80, Jan. 2009.
- M. Defferrard, X. Bresson, and P. Vandergheynst, "Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering," in *Advances in Neural Information Processing Systems 29*, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, Eds. Curran Associates, Inc., 2016, pp. 3844–3852.
- P. W. Battaglia et al., "Relational inductive biases, deep learning, and graph networks," *arXiv [cs.LG]*, 04-Jun-2018.
- M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst, "Geometric Deep Learning: Going beyond Euclidean data," *IEEE Signal Process. Mag.*, vol. 34, no. 4, pp. 18–42, Jul. 2017.
- W. L. Hamilton, R. Ying, and J. Leskovec, "Representation Learning on Graphs: Methods and Applications," *arXiv [cs.SI]*, 17-Sep-2017.
- Z. Zhang, P. Cui, and W. Zhu, "Deep Learning on Graphs: A Survey," *arXiv [cs.LG]*, 11-Dec-2018.
- Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A Comprehensive Survey on Graph Neural Networks," *arXiv [cs.LG]*, 03-Jan-2019.
- J. Zhou et al.**, "Graph Neural Networks: A Review of Methods and Applications," *arXiv [cs.LG]*, 20-Dec-2018.
- Y. Li, D. Tarlow, M. Brockschmidt, and R. Zemel, "Gated Graph Sequence Neural Networks," *arXiv [cs.LG]*, 17-Nov-2015.
- L. Ruiz, F. Gama, and A. Ribeiro, "Gated Graph Convolutional Recurrent Neural Networks," *arXiv [cs.LG]*, 05-Mar-2019.
- H. Dai, Z. Kozareva, B. Dai, A. Smola, and L. Song, "Learning Steady-States of Iterative Algorithms over Graphs," in *Proceedings of the 35th International Conference on Machine Learning*, 2018, vol. 80, pp. 1106–1114.
- J. You, R. Ying, X. Ren, W. L. Hamilton, and J. Leskovec, "GraphRNN: Generating Realistic Graphs with Deep Auto-regressive Models," *arXiv [cs.LG]*, 24-Feb-2018.
- J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun, "Spectral Networks and Locally Connected Networks on Graphs," *arXiv [cs.LG]*, 21-Dec-2013.
- O. Levy and Y. Goldberg, "Neural Word Embedding as Implicit Matrix Factorization," in *Advances in Neural Information Processing Systems 27*, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2014, pp. 2177–2185.
- D. K. Duvenaud et al., "Convolutional Networks on Graphs for Learning Molecular Fingerprints," in *Advances in Neural Information Processing Systems 28*, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, Eds. Curran Associates, Inc., 2015, pp. 2224–2232.
- P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph Attention Networks," *arXiv [stat.ML]*, 30-Oct-2017.
- J. Zhang, X. Shi, J. Xie, H. Ma, I. King, and D.-Y. Yeung, "GaAN: Gated Attention Networks for Learning on Large and Spatiotemporal Graphs," *arXiv [cs.LG]*, 20-Mar-2018.
- D. Wang, P. Cui, and W. Zhu, "Structural Deep Network Embedding," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, San Francisco, California, USA, 2016, pp. 1225–1234.
- T. N. Kipf and M. Welling, "Variational Graph Auto-Encoders," *arXiv [stat.ML]*, 21-Nov-2016.
- Y. Li, D. Vinyals, C. Dyer, R. Pascanu, and P. Battaglia, "Learning Deep Generative Models of Graphs," *arXiv [cs.LG]*, 08-Mar-2018.
- M. Simonovsky and N. Komodakis, "GraphVAE: Towards Generation of Small Graphs Using Variational Autoencoders," in *Artificial Neural Networks and Machine Learning – ICANN 2018*, 2018, pp. 412–422.



- N. De Cao and T. Kipf, "MolGAN: An implicit generative model for small molecular graphs," arXiv [stat.ML], 30-May-2018.
- S. Pan, R. Hu, S.-F. Fung, G. Long, J. Jiang, and C. Zhang, "Learning Graph Embedding With Adversarial Training Methods," IEEE Transactions on Cybernetics. pp. 1-13, 2019.
- A. Sperduti and A. Starita, "Supervised neural networks for the classification of structures," IEEE Trans. Neural Netw., vol. 8, no. 3, pp. 714-735, 1997.
- M. Gori, G. Monfardini, and F. Scarselli, "A new model for learning in graph domains," in Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005., 2005, vol. 2, pp. 729-734 vol. 2.
- M. A. Khamsi and W. A. Kirk, An Introduction to Metric Spaces and Fixed Point Theory. John Wiley & Sons, 2011.
- J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, "Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling," arXiv [cs.NE], 11-Dec-2014.
- C. Gallicchio and A. Micheli, "Graph Echo State Networks," in The 2010 International Joint Conference on Neural Networks (IJCNN), 2010, pp. 1-8.
- C. Zhuang and Q. Ma, "Dual graph convolutional networks for graph-based semi-supervised classification," Proceedings of the 2018 World Wide Web Conference, 2018.
- D. Bahdanau, K. Cho, and Y. Bengio, "Neural Machine Translation by Jointly Learning to Align and Translate," arXiv [cs.CL], 01-Sep-2014.
- A. Vaswani et al., "Attention is All you Need," in Advances in Neural Information Processing Systems 30, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds. Curran Associates, Inc., 2017, pp. 5998-6008.
- J. B. Lee, R. A. Rossi, S. Kim, N. K. Ahmed, and E. Koh, "Attention Models in Graphs: A Survey," arXiv [cs.AI], 20-Jul-2018.
- P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P.-A. Manzagol, "Stacked Denoising Autoencoders: Learning Useful Representations in a Deep Network with a Local Denoising Criterion," J. Mach. Learn. Res., vol. 11, no. Dec, pp. 3371-3408, 2010.
- D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," arXiv [stat.ML], 20-Dec-2013.
- I. Goodfellow et al., "Generative Adversarial Nets," in Advances in Neural Information Processing Systems 27, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2014, pp. 2672-2680.
- X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local neural networks," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 7794-7803.
- 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, 2017, pp. 1263-1272.
- A. Paszke et al., "Automatic differentiation in PyTorch," 28-Oct-2017.
- M. Abadi et al., "Tensorflow: A system for large-scale machine learning," in 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI'16), 2016, pp. 265-283.
- M. Wang et al., "Deep graph library, 2018," URL <http://dgl.ai>.
- Alibaba, "Euler," 2019.
- M. Fey and J. E. Lenssen, "Fast Graph Representation Learning with PyTorch Geometric," arXiv [cs.LG], 06-Mar- 2019.
- D. Chai, L. Wang, and Q. Yang, "Bike Flow Prediction with Multi-graph Convolutional Networks," in Proceedings of the 26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, Seattle, Washington, 2018, pp. 397-400.
- M. Wang et al., "Dynamic Spatio-temporal Graph-based CNNs for Traffic Prediction," arXiv [cs.CV], 05-Dec- 2018.
- Y. Li, R. Yu, C. Shahabi, and Y. Liu, "Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting," arXiv [cs.LG], 06-Jul-2017.
- Z. Wu, S. Pan, G. Long, J. Jiang, and C. Zhang, "Graph WaveNet for Deep Spatial-Temporal Graph Modeling," arXiv [cs.LG], 31-May-2019.
- Y. Zhang, T. Cheng, and Y. Ren, "A graph deep learning method for short-term traffic forecasting on large road networks," Computer-Aided Civil and Infrastructure Engineering. 2019.
- F. L. Opolka, A. Solomon, C. Cangea, P. Veličković, P. Liò, and R. Devon Hjelm, "Spatio-Temporal Deep Graph Infomax," arXiv [cs.LG], 12-Apr-2019.



- C. Chen et al., "Gated Residual Recurrent Graph Neural Networks for Traffic Prediction," in Proceedings of the AAAI Conference on Artificial Intelligence, 2019, vol. 33, pp. 485-492.
- N. Zhang, X. Guan, J. Cao, X. Wang, and H. Wu, "A Hybrid Traffic Speed Forecasting Approach Integrating Wavelet Transform and Motif-based Graph Convolutional Recurrent Neural Network," arXiv [cs.CV], 14-Apr- 2019.
- X. Cheng, R. Zhang, J. Zhou, and W. Xu, "DeepTransport: Learning Spatial-Temporal Dependency for Traffic Condition Forecasting," in 2018 International Joint Conference on Neural Networks (IJCNN), 2018, pp. 1-8.
- J. Sun, J. Zhang, Q. Li, X. Yi, and Y. Zheng, "Predicting Citywide Crowd Flows in Irregular Regions Using Multi-View Graph Convolutional Networks," arXiv [cs.CV], 19-Mar-2019.
- B. Wang, X. Luo, F. Zhang, B. Yuan, A. L. Bertozzi, and P. Jeffrey Brantingham, "Graph-Based Deep Modeling and Real Time Forecasting of Sparse Spatio-Temporal Data," arXiv [cs.LG], 02-Apr-2018.
- B. Yu, H. Yin, and Z. Zhu, "Spatio-Temporal Graph Convolutional Networks: A Deep Learning Framework for Traffic Forecasting," arXiv [cs.LG], 14-Sep-2017.
- B. Yu, H. Yin, and Z. Zhu, "ST-UNet: A Spatio-Temporal U-Net for Graph-structured Time Series Modeling," arXiv [cs.LG], 13-Mar-2019.
- L. Zhao et al., "T-GCN: A Temporal Graph Convolutional Network for Traffic Prediction," arXiv [cs.LG], 12- Nov-2018.
- H. Yao, X. Tang, H. Wei, G. Zheng, and Z. Li, "Revisiting spatial-temporal similarity: A deep learning framework for traffic prediction," in AAAI Conference on Artificial Intelligence, 2019.
- J. Atwood and D. Towsley, "Diffusion-convolutional neural networks," in Advances in Neural Information Processing Systems, 2016, pp. 1993-2001.
- X. Zhou, Y. Shen, and L. Huang, "Revisiting Flow Information for Traffic Prediction," arXiv [eess.SP], 03-Jun- 2019.
- Y. Li and C. Shahabi, "A Brief Overview of Machine Learning Methods for Short-term Traffic Forecasting and Future Directions," SIGSPATIAL Special, vol. 10, no. 1, pp. 3-9, Jun. 2018.

