

Dynamic Graph Echo State Networks

Domenico Tortorella¹ and Alessio Micheli¹

1- University of Pisa - Department of Computer Science
Largo B. Pontecorvo 3, 56127 Pisa - Italy

Abstract. Dynamic temporal graphs represent evolving relations between entities, e.g. interactions between social network users or infection spreading. We propose an extension of graph echo state networks for the efficient processing of dynamic temporal graphs, with a sufficient condition for their echo state property, and an experimental analysis of reservoir layout impact. Compared to temporal graph kernels that need to hold the entire history of vertex interactions, our model provides a vector encoding for the dynamic graph that is updated at each time-step without requiring training. Experiments show accuracy comparable to approximate temporal graph kernels on twelve dissemination process classification tasks.

1 Introduction

Graphs are relevant in modelling entities and relations between them, e.g. atoms and bonds in a molecule, or paper citation networks. A plethora of machine learning models able to treat directly these data structures has been proposed and successfully employed for classification and regression tasks, both on graphs and on vertices [1]. However, some relations evolve through time. People spreading a disease by interacting at certain instants, or users sharing posts on a social network, are just two examples of *dissemination processes* (Fig. 1). Many approaches have been recently proposed for learning representation of dynamic graphs [2]. In particular, temporal graph kernels [3] extend classic kernels employed for static graphs by transforming dynamic graphs into equivalent static ones; spatio-temporal graph convolutional networks [4] deal directly with dynamic graphs, but require end-to-end training. On the other hand, efficient reservoir approaches such as echo state networks (ESN), which embed data structures without requiring recurrent weights training, have been successfully applied to classification and regression tasks on static graphs [5, 6].

In section 2 we present DynGESN¹, an extension of GraphESN for dynamic temporal graphs, along with a sufficient condition for the echo state property. In section 3 we evaluate our model on twelve classification tasks, analysing the impact of reservoir layout on accuracy and comparing performances with temporal graph kernels.

2 Model

We define a *dynamic* graph \mathcal{G} as a pair $(\mathcal{V}, \mathcal{E})$, where \mathcal{V} is the set of vertices, and $\mathcal{E} = \{(u, v, t) \mid u, v \in \mathcal{V}, t \in 1..T\}$ is the set of edges $\textcircled{u} \rightarrow \textcircled{v}$ between a pair

¹Code available at github.com/dtortorella/dyngraphesn, based on the original GraphESN implementation by C. Gallicchio.




Fig. 1: The spreading of an infection, an example of dissemination process. Infected vertices (in red, $\mathbf{u}_t(v) = 1$) contaminate susceptible vertices (in black, $\mathbf{u}_t(v) = 0$) by interacting with them at time-step t (temporal edges).

of vertices at a time-step t . The graph \mathcal{G} can be characterized as *undirected* if $(u, v, t) \in \mathcal{E} \Rightarrow (v, u, t) \in \mathcal{E}$, and as *static* if $(u, v, t) \in \mathcal{E} \Rightarrow (u, v, t') \in \mathcal{E} \forall t, t'$ (i.e., constant edge set). We also define $\mathcal{E}_t = \{(u, v) \mid (u, v, t) \in \mathcal{E}\}$, the neighbourhood $\mathcal{N}_t(v) = \{u \in \mathcal{V} \mid (u, v) \in \mathcal{E}_t\}$, and \mathbf{A}_t as the adjacency matrix at time-step t . Finally, we associate at each vertex v a label sequence $\mathbf{u}_t(v) \in \mathbb{R}^U$ (e.g. $\mathbf{u}_t(v) \in \{0, 1\}$ in Fig. 1).

The ESNs are a particular class of recurrent neural networks (RNN) in which the recurrent weights are randomly initialized under certain conditions (see later) and kept fixed, while only a memoryless readout layer is trained. Formally, an ESN is characterized as an input-driven dynamical system governed by a transition function $\mathbf{x}_t = F(\mathbf{u}_t, \mathbf{x}_{t-1})$ with states \mathbf{x}_t belonging to a compact subset \mathcal{X} of \mathbb{R}^H , and $\mathbf{u}_t \in \mathbb{R}^U$ being the input at time t . An input sequence $\mathbf{u} = [\mathbf{u}_1 \dots \mathbf{u}_T]$ is encoded into the embedding space \mathcal{X} by applying iteratively F from $t = 1$ to the end, with $\mathbf{x}_0 \in \mathcal{X}$ as initial state for $t = 0$.

We extend the deep ESN reservoir model for vector sequences [7] to dynamic temporal graphs by having each recurrent layer $\ell \in 1..L$ compute new vertex features $\mathbf{x}_t^{(\ell)}(v) \in \mathcal{X} = [-1, +1]^H$ by the transition function $F^{(\ell)}(\mathbf{u}_t, \mathbf{A}_t, \mathbf{x}_{t-1})$ defined vertex-wise as

$$\mathbf{x}_t^{(\ell)}(v) = \gamma_\ell \tanh \left(\mathbf{W}_{\text{in}}^{(\ell)} \mathbf{x}_t^{(\ell-1)}(v) + \sum_{v' \in \mathcal{N}_t(v)} \hat{\mathbf{W}}^{(\ell)} \mathbf{x}_{t-1}^{(\ell)}(v') \right) + (1 - \gamma_\ell) \mathbf{x}_{t-1}^{(\ell)}(v) \quad (1)$$

with $0 < \gamma_\ell \leq 1$ as leakage constant, $\mathbf{x}_t^{(0)}(v) = \mathbf{u}_t(v)$ for layer $\ell = 1$, and initial state $\mathbf{x}_0^{(\ell)}(v) = \mathbf{0}$. Both input weights $\mathbf{W}_{\text{in}}^{(\ell)}$ and recurrent weights $\hat{\mathbf{W}}^{(\ell)}$ are randomly initialized. The embedding of a graph \mathcal{G} is then given by the pooled vertex features of the final state for each layer

$$\mathbf{X}_{\mathcal{G}} = \begin{bmatrix} \sum_{v \in \mathcal{V}} \mathbf{x}_T^{(1)}(v) & \dots & \sum_{v \in \mathcal{V}} \mathbf{x}_T^{(L)}(v) \end{bmatrix} \in \mathbb{R}^{HL} \quad (2)$$

and can be used to perform regression or classification, e.g. by training a linear readout

$$\mathbf{y} = \mathbf{W}_{\text{out}} \mathbf{X}_{\mathcal{G}} + \mathbf{b} \quad (3)$$

by ridge regression, or by training a support vector machine (SVM).

Equation (1) resembles closely the transition function of GraphESN [5], except for the use of temporal neighbourhood \mathcal{N}_t due to the change in graph connectivity between time-steps. However, they operate rather differently: GraphESN processes static graphs with constant vertex labels by iterating its transition function until the global state converges to a fixed point (t tracking iterations in this case), while ours acts analogously to ESN for finite sequences (in our case a sequence of T graphs over time t). Furthermore, the deep GraphESN [6] waits for convergence in the previous layer to pass the final state as input to the subsequent layer, while our model works more similarly to deep ESN for sequences, i.e. passing a state to the subsequent layer at each time-step [7].

ESN are able to provide meaningful encoding of sequences thanks to the echo state property (ESP) [7], which ensures that perturbations in the initial state are ‘washed out’ in the long term. Furthermore, the contractivity of F , i.e. $\|F(\mathbf{u}, \mathbf{x}) - F(\mathbf{u}, \mathbf{x}')\| \leq C \|\mathbf{x} - \mathbf{x}'\|$, $C < 1$ (reducing embedding state distance under same input), ensures that the embedding space of sequences has a suffix-based organization [8].

Proposition 1 (DynGESN ESP). The transition function defined in equation (1) is contractive with constant $C_t = (1 - \gamma^{(\ell)}) + \gamma^{(\ell)} \|\hat{\mathbf{W}}^{(\ell)}\| \|\mathbf{A}_t\|$ for each layer ℓ . Therefore, a sufficient condition for the ESP to hold is

$$\|\hat{\mathbf{W}}^{(\ell)}\| < \frac{1}{\lceil \alpha_t^* \rceil} \quad \text{for all layers } \ell \in 1..L, \quad (4)$$

where $\alpha_t^* = \|\mathbf{A}_t\|$ is the maximum eigenvalue/singular value of the adjacency matrix \mathbf{A}_t , and $\lceil \alpha_t^* \rceil$ is the geometric mean $\lceil \alpha_t^* \rceil \stackrel{\text{def}}{=} \limsup_{T \rightarrow \infty} \sqrt[T]{\prod_{1 \leq t \leq T} \alpha_t^*}$.

The proof works out analogously to the sufficient condition for ESN in [7], and is omitted for brevity. Interestingly, Proposition 1 provides a weaker ESP sufficient condition for GraphESN [6]: in case of static graphs $\lceil \alpha_t^* \rceil = \alpha^* \leq k^*$, where k^* is the maximum vertex degree, and $\alpha^* = k^*$ for complete graphs.

3 Experiments and discussion

We evaluate our model for dynamic graph binary classification on twelve datasets developed by Oettershagen et al. [3], where dissemination processes based on the *susceptible-infected* (SI) epidemic model have been simulated on six different real-world social interaction datasets. In a SI model vertices are labelled either *susceptible* or *infected*, switching label from former to latter with fixed probability p at each time-step when they are directly linked to an infected vertex.

The two classes of the first six tasks (ct1) are dynamic graphs that follow a SI process with probability $p = 0.5$, and dynamic graphs whose vertices randomly switch from susceptible to infected. In the second six tasks (ct2) classifiers have to discriminate between two SI dissemination processes with contagion probability $p = 0.2$ and $p = 0.8$. We refer to [3] for further details.

Since our model depends on random weights initialization, we evaluated classification accuracies by averaging on 200 bootstraps with 90%-10% training/test

splits instead of a 10-fold cross-validation. Reservoir weights are randomly generated following an uniform distribution, with recurrent weights rescaled to $\|\hat{\mathbf{W}}^{(\ell)}\| = 0.9 \alpha_{\text{mean}}$ in order to satisfy Proposition 1, α_{mean} being the average $[\alpha_t^*]$ on each dataset; we fixed $\gamma_\ell = 0.1$. A linear readout (3) is trained by ridge regression with regularization $\lambda = 10^{-3}$.

Figure 2 shows how both hidden state dimension (i.e. reservoir size per layer) $H \in \{1, 2, 4, 8, 16\}$ and number of layers $L \in 1..6$ affect classification accuracy. We notice two trends:

- (i) increasing the number of hidden units H per layer with fixed depth L increases accuracy, which can be explained by the ability of larger reservoirs to offer richer dynamics;
- (ii) increasing the number of layers with fixed H has overall a significant impact on accuracy, with a steep increase up to $L = 3$; the beneficial effect of depth has already been reported for GraphESN [6], and in our case could be explained similarly to deep ESN for sequences [9], i.e. that deeper layers represent different time-scales of a graph sequence.




Fig. 2: Impact of reservoir layout on classification accuracy.

We now compare DynGESN against other models. Temporal graph kernels compute the Gram matrix for a SVM by first transforming the dynamic temporal graphs into static, and then by applying ‘classic’ graph kernels such as k -step random walk (RW- k) or Weisfeiler–Lehman sub-tree kernel of depth h (WL- h). In [3] three static transformations are proposed: *reduced graph representation* (RD), which loses temporal information, but does not increase graph dimension; *direct line expansion* (DL) and *static expansion* (SE), which preserve temporal information by providing a much larger static graph representation (Table 3). Temporal graph kernels APPR- S are also proposed to approximate the random-

walk kernel on direct line expansion (i.e. DL-RW) by sampling k -step random walks starting on S vertices of the expanded graph. Since these offer the best trade-off between speed and accuracy on temporal graph kernels, we will consider them in our comparison.

Tables 1 and 2 compare the classification accuracy (with standard deviation, computed on bootstraps) of our model for $H = 16$ and $L = 4$ against the results reported in [3] for temporal graph kernels. Best accuracies between DynGESN and APPR- S are highlighted in bold; results for exact DL temporal graph kernels are also reported for reference, with SE having similar performances to DL, and RD significantly poorer with respect to both.

DynGESN performs consistently better than APPR-100, and on par or above APPR-250 on all datasets except from mit-ct1 and mit-ct2, which are problematic also for some exact temporal kernel due to their small number of samples. Notice also that S must scale with dataset size in order to satisfy the theoretical approximation guarantees for APPR- S kernels [3].

Model	dblp	facebook	highschool	infectious	mit	tumblr
DL-RW	98.7 \pm 0.1	96.5 \pm 0.1	97.4 \pm 0.7	98.0 \pm 0.4	92.9 \pm 0.9	95.2 \pm 0.6
DL-WL	98.5 \pm 0.2	96.6 \pm 0.2	99.2 \pm 0.6	98.1 \pm 0.4	91.7 \pm 1.6	94.2 \pm 0.4
APPR-50	93.1 \pm 0.5	89.4 \pm 0.5	87.6 \pm 1.7	83.4 \pm 1.6	82.8 \pm 2.0	89.3 \pm 0.7
APPR-100	95.7 \pm 0.5	92.4 \pm 0.4	90.2 \pm 1.9	91.1 \pm 1.0	85.1 \pm 2.6	90.4 \pm 1.0
APPR-250	97.2 \pm 0.2	94.6 \pm 0.3	94.0 \pm 1.3	95.1 \pm 0.9	89.3 \pm 2.8	92.7 \pm 0.3
DynGESN	97.7 \pm 1.8	93.0 \pm 2.5	94.4 \pm 5.3	94.7 \pm 5.3	69.7 \pm 9.2	93.3 \pm 3.9

Table 1: Classification accuracy on ct1 tasks.

Model	dblp	facebook	highschool	infectious	mit	tumblr
DL-RW	81.8 \pm 0.9	80.0 \pm 0.5	93.4 \pm 1.0	88.7 \pm 1.2	82.6 \pm 2.1	77.2 \pm 1.0
DL-WL	76.5 \pm 1.0	80.0 \pm 0.5	89.3 \pm 0.7	78.7 \pm 1.5	36.4 \pm 4.0	78.2 \pm 1.3
APPR-50	69.7 \pm 0.8	72.0 \pm 0.7	76.4 \pm 3.0	74.6 \pm 1.8	60.4 \pm 4.0	74.7 \pm 1.3
APPR-100	74.5 \pm 0.7	73.1 \pm 0.6	83.8 \pm 1.8	75.6 \pm 1.9	65.4 \pm 4.2	76.5 \pm 1.5
APPR-250	76.4 \pm 0.9	77.9 \pm 0.6	90.4 \pm 1.8	78.6 \pm 2.1	66.9 \pm 2.5	78.4 \pm 1.3
DynGESN	74.3 \pm 4.7	76.1 \pm 3.9	92.8 \pm 5.2	80.6 \pm 9.1	63.3 \pm 11.0	76.8 \pm 6.2

Table 2: Classification accuracy on ct2 tasks.

DynGESN is able to obtain an embedding for each time-step needing only previous vertex features and current labels and edges, thus requiring $O(V)$ space and $O(V + E_t)$ matrix operation at each time-step ($O(TV + E)$ overall; $\sim 10^{-4}s$ per time-step on a laptop for our implementation). Thus our model can be applied in an ‘on-line’ setting, updating vertex and graph embedding directly as the dynamic graph evolves, without needing anything more than \mathbf{x}_{t-1} , \mathbf{u}_t and \mathcal{E}_t at each time-step. Furthermore, DynGESN can be easily adopted for vertex regression/classification tasks by simply skipping graph pooling (2).

PROCESSING GRAPH SIZE			EMBEDDING TIME COMPLEXITY			
	vertices	edges		RW- k	APPR- S	WL- h
RD	$O(V)$	$O(V^2)$	RD	$O(V^{2k})$	—	$O(N^2 h V^2)$
DL	$O(E)$	$O(E^2)$	DL	$O(E^{2k})$	$O(Sk)$	$O(NhE^2 + N^2 hE)$
SE	$O(E)$	$O(E)$	SE	$O(E^k)$	—	$O(N^2 hE)$
DynGESN	$O(V)$	$O(E_t)$	DynGESN	$O(V + E_t)$ per time-step		

Table 3: Comparison of space requirements and time complexity (V , E and E_t are the cardinalities of \mathcal{V} , \mathcal{E} and \mathcal{E}_t , respectively, for an input dynamic graph; N is number of dataset samples).

4 Conclusions

We have presented an extension of static graph echo state networks for the efficient processing of dynamic graphs with time-dependent vertex labels, along with theoretical conditions for the echo state property to hold. Experiments performed on twelve dissemination process classification tasks show an accuracy comparable to approximate temporal graph kernels paired with SVMs. These results, along with significantly lower space requirements, the ability to embed graphs *on-line* (thus to amortise the cost on time-steps), and the lack of training for the embedding function, suggest that our method is the most promising for scaling on larger real-world applications. Finally, our model can also be easily employed for classification/regression tasks on vertices, which will be examined in subsequent works.

References

- [1] D. Bacciu, F. Errica, A. Micheli, and M. Podda. A gentle introduction to deep learning for graphs. *Neural Networks*, 129:203–221, 2020.
- [2] S. M. Kazemi, R. Goel, K. Jain, I. Kobyzev, A. Sethi, P. Forsyth, and P. Poupart. Representation learning for dynamic graphs: A survey. *Journal of Machine Learning Research*, 21(70):1–73, 2020.
- [3] L. Oettershagen, N. M. Kriege, C. Morris, and P. Mutzel. Temporal graph kernels for classifying dissemination processes. In *Proceedings of the 2020 SIAM International Conference on Data Mining*, pages 496–504, 2020.
- [4] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu. A Comprehensive Survey on Graph Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32(1):4–24, 2021.
- [5] C. Gallicchio and A. Micheli. Graph echo state networks. In *The 2010 International Joint Conference on Neural Networks (IJCNN)*, pages 3967–3974. IEEE, July 2010.
- [6] C. Gallicchio and A. Micheli. Fast and Deep Graph Neural Networks. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI-20)*, November 2020.
- [7] C. Gallicchio and A. Micheli. Echo state property of deep reservoir computing networks. *Cognitive Computation*, 9(3):337–350, 2017.
- [8] C. Gallicchio and A. Micheli. Architectural and markovian factors of echo state networks. *Neural Networks*, 24(5):440–456, 2011.
- [9] C. Gallicchio, A. Micheli, and L. Pedrelli. Deep reservoir computing: A critical experimental analysis. *Neurocomputing*, 268:87–99, 2017.