

When to Remember Where You Came from: Node Representation Learning in Higher-order Networks

Caleb Belth, Fahad Kamran, Donna Tjandra, Danai Koutra
University of Michigan

Email: {cbelth, fhdkmrr, dotjandr, dkoutra}@umich.edu

Abstract—For trajectory data that tend to have beyond first-order (i.e., non-Markovian) dependencies, higher-order networks have been shown to accurately capture details lost with the standard aggregate network representation. At the same time, representation learning has shown success on a wide range of network tasks, removing the need to hand-craft features for these tasks. In this work, we propose a node representation learning framework called EVO or *Embedding Variable Orders*, which captures non-Markovian dependencies by combining work on higher-order networks with work on node embeddings. We show that EVO outperforms baselines in tasks where high-order dependencies are likely to matter, demonstrating the benefits of considering high-order dependencies in node embeddings. We also provide insights into when it does or does not help to capture these dependencies. To the best of our knowledge, this is the first work on representation learning for higher-order networks.

I. INTRODUCTION

Recent work on higher-order networks¹ (HONs) [2], [3] has demonstrated the importance of considering non-Markovian dependencies when building a network representation from trajectory data (e.g., career paths, flight or ship itineraries, clickstreams, etc. [2], [3], [4]). Meanwhile, representation learning has been useful for learning feature representations for standard (first-order) networks that can be directly used for downstream tasks [5]. While high-order dependencies are important for accurately *modeling* a network, it is not immediately obvious that the added modeling accuracy is *always* useful to learn over. To understand when it is useful, we evaluate the performance of learning over HONs on several classification and clustering tasks, and compare to the performance of learning over standard network representations.

We focus on two research questions: (RQ1) how can we capture non-Markovian dependencies in feature representations and (RQ2) when are these dependencies useful? In response to the first question, we propose EVO or “Embedding Variable Orders” (Figure 1), a modular framework for representation learning that works with HON representations to

¹Note that we are using “higher-order” to refer to *non-Markovian* or *beyond first-order* dependencies in trajectories, rather than motifs [1].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ASONAM '19, August 27-30, 2019, Vancouver, Canada

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6868-1/19/08

<http://dx.doi.org/10.1145/3341161.3342911>




Figure 1: The proposed EVO framework (white background), contrasted with the standard approach SWNE (gray background) (§ IV-B). EVO takes trajectories as input (e.g., flight itineraries). Step 1 creates a HON that may contain multiple versions of a node (e.g., DTW is mapped to two conditional nodes (DTW|LGA & DTW|DSM) since the 0.5 probability of visiting SFO from DTW becomes 0.67 or 0.33 depending on how DTW was reached). Step 2 learns node embeddings from the HON. Step 3 combines the embeddings of all node versions into a single embedding capturing the important properties of *all* its versions. SWNE has only two steps.

learn dependency-preserving node representations. In response to the second question, we find that EVO is effective at capturing high-order dependencies and is useful for a range of tasks related to network trajectories. However, contrary to what might be expected, we find that capturing high-order dependencies for tasks where they are not very relevant is not only ineffective, but can be *detrimental*, suggesting that capturing more information is not always helpful. This is a significant finding, since HON representations consist of multiple nodes for multiple dependencies, requiring additional computational power—something highly undesirable for large graphs. Our main contributions are:

- We are the first to propose a **general, modular framework**, combining work in representation learning and HONs to generate node embeddings that capture non-Markovian dependencies. It readily supports both neighborhood and structural embeddings, and works with any embedding approach.
- Through diverse experiments, we bring insights into **when downstream tasks can benefit** from these embeddings.

II. RELATED WORK

Higher-order Networks attempt to preserve non-first-order dependencies in a network representation. While works such as [6] used a fixed-order representation, usually only modeling second-order dependencies, recent work focused on variable-order dependencies, which are automatically identified from trajectories [2], [3]. Xu *et al.* [2] build a variable-order network by creating multiple versions of nodes to capture

TABLE I: Major symbols and their meanings.

V	set of nodes	$G(V, E)$	standard network
T	trajectories over nodes V	$H(V_H, E_H, \delta)$	HON w/ mapping from V to V_H
\mathbf{r}_v	embedding of node v		

multiple dependencies. The resulting HON admits standard graph methods. Scholtes *et al.* [3] build a layered network (which may require tailored graph methods), with each layer capturing a progressively higher dependency level.

Representation Learning aims to learn feature representations for nodes (or other structures) which can be directly used for network tasks. Many methods are inspired by word embeddings [7], replacing the linear context around words with graph structure around nodes by using variants of random walks. Some are proximity-based, e.g. [8], [9], while others are structural, e.g. [10]. Other methods, e.g. [11], utilize the connection between matrix factorization and skip-gram with negative sampling [12]. For more, [5] surveys the extensive work in representation learning. To the best of our knowledge, we are the first to look at representation learning for HONs. While [1] uses similar terminology, it is not closely related, using “higher-order networks” to refer to graph motifs, whereas we use it to refer to non-Markovian dependencies.

III. METHOD

A. Preliminary Definitions

A (standard or first-order) *network* or *graph* $G = (V, E)$ is a set of nodes and a set of edges among them, optionally with weights denoting the connection strength. G implicitly encodes the Markov assumption.

We follow [2] and define a *higher-order network (HON)* $H = (V_H, E_H, \delta)$ as a network, along with a bijective function $\delta : V \rightarrow \mathcal{P}(V_H)$ mapping each node $v \in V$ to the subset of nodes in V_H encoding v ’s dependencies. For instance, in Figure 1, $\delta(\text{DTW}) = \{\text{DTW}|\text{LGA}, \text{DTW}|\text{DSM}\}$. Its inverse maps the nodes back: $\delta^{-1}(\{\text{DTW}|\text{LGA}, \text{DTW}|\text{DSM}\}) = \text{DTW}$.

B. Problem Definition

Problem 1. *Given a set of trajectories T over nodes V , where the Markov assumption may be violated, find a representation $\mathbf{r}_v \in \mathbb{R}^d$ for each node $v \in V$ such that the representations capture the non-Markovian dependencies.*

C. EVO: Embedding Variable Orders

Our proposed method, EVO, works in three steps:

(S1) Construct Higher-order Network. Any approach for building a HON H (e.g. [2]) can be used if there exists a function δ and its inverse δ^{-1} to map nodes to and from V_H .

(S2) Learn Representations from HON. For each node $h \in V_H$, we obtain its representation \mathbf{r}_h . Assuming that the HON is a conventional graph (e.g., not layered), any representation learning method can be used [5].

(S3) Combine Representations. This step uses a function $f : \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}^d$ to create a single representation for each node in V from its corresponding nodes in V_H . For any node v , we get its embedding (denoted \mathbf{r}_v) by applying f (e.g., average, max, etc.) to the embeddings of all the nodes in V_H which

TABLE II: Data statistics: median length of trajectories, max dependency found by [2], size of G (in SNE form), and size of H .

Dataset	Med. Len.	Max Dep.	$ V $	$ E $	$ V_H $	$ E_H $
Us Flights [3]	5	4	175	1,598	9,776	49,700
London Tube [3]	12	2	268	646	1,029	2,073

correspond to v (which may vary in number per node). Let $R = \{\mathbf{r}_h : h \in \delta(v)\}$ be the set of embeddings of nodes in $\delta(v)$, acquired in **S2**. Formally, $\mathbf{r}_v = f(R) \forall v \in V$.

In summary, the EVO framework allows for a number of methods, each requiring the selection of (1) a HON representation, (2) a representation learning method, and (3) a function f to convert the representations of HON nodes V_H into representations for nodes V . Its runtime depends on the methods employed in **S1** and **S2**. Assuming that their complexities are α and β respectively, and given that **S3** can be performed with a hashtable in a pass over V_H , in $\mathcal{O}(|V_H|)$ time, EVO’s total runtime is $\mathcal{O}(\alpha + \beta + |V_H|)$.

IV. EMPIRICAL ANALYSIS

In our experimental analysis, we aim to answer two main research questions: **[RQ1] Performance:** Does EVO effectively capture high-order dependencies? **[RQ2] Appropriate Uses:** In what situations does it make sense to capture high-order dependencies? We consider a variety of node classification and clustering tasks.

A. Datasets

We use two datasets from [3] (Table II): (1) A London Tube dataset containing passenger itineraries from the London subway system and (2) a US Flights dataset containing flight itineraries, such as LGA \rightarrow DTW \rightarrow SFO (Figure 1), meaning someone flew from LaGuardia to Detroit to San Francisco.

B. Baselines

We consider three baselines:

- **SNE** (“Standard Network Embedding”) aggregates the trajectories into a standard network representation, where an edge exists between two nodes if the edge occurred in at least one trajectory. Node representations are learned via the same embedding approach used in EVO.
- **SWNE** (“Standard Weighted Network Embedding”) is SNE, but introduces edge weights to count edge occurrences.
- **ARWE** (“Altered-Random-Walk Embedding”) is an approach we introduce. ARWE runs variable-order Markov random walks to generate context for skip-gram-based representation learning directly, *without* building the intermediate HON. The random walks simulate trajectories by obeying the same dependencies used to build H .

C. Experimental Setup

For **S1** of EVO, we construct a HON using [2] without using the optional min-support parameter. For **S2**, we use node2vec [9], struc2vec [10], or xNetMF [11], depending on the task (using node2vec unless the task is structural). xNetMF and struc2vec do not use weights. Furthermore, xNetMF does not use random walks, while struc2vec’s walks are on a k -level graph, so SWNE and ARWE do not apply. For node2vec

and `struc2vec`, we set the walk length to be the median trajectory length for that dataset (e.g. 5 for Flights), to treat random walks as simulations of real trajectories. We perform 50 walks per node, and use a window-size of 10. We use none of `struc2vec`’s computational optimizations. For **S3**, we consider multiple options—max, min, average, and Hadamard product—for f , each of which is applied elementwise. We also tried concatenation, but saw similar or worse results. We found that max and average worked best, so we report results for only these, named **EVO-max** and **EVO-avg**.

Results are reported as averages with standard deviations over 10 runs. The top result for each metric is shaded gray. If the result is statistically significant (which we consider $p \leq 0.05$ for a paired t-test), we mark it with an “*”. We discuss results in each task’s section, and takeaways in § IV-F.

D. Node Classification

We designed four binary classification tasks to answer our research questions. We use node representations of each method as input to a random forest classifier with 100 trees, running with 10 different 70%/30% train/test splits (the same splits for each method). When using `node2vec`, we tune its p and q parameters as in [9], using grid search with 10-fold cross-validation. We report Accuracy (ACC), AUROC (AUC), and F1 score in Table III, and discuss results in § IV-F.

1) Trajectory Node Classification is designed to investigate EVO’s usefulness for tasks related to trajectories. In this task, we attempt to predict whether each London Tube station services one line or multiple lines. Since trains follow lines, this task is related to trajectories, requiring memory of where a train came from to distinguish between traffic that comes from one line vs. traffic that comes from multiple lines.

Setup. Since the task is not structural, we use our framework with `node2vec`. Note that 70.9% of stations service one line.

Results. EVO’s benefit is clear, as it outperforms all baselines by a statistically significant amount by all metrics, demonstrating that EVO captures trajectories and can identify traffic from multiple lines. In contrast, the baselines only capture *how much* traffic comes through each station. To help explain this, we observe that stations with multiple lines have 6.12 dependencies on average, while stations with only one line have only 2.91. Thus, when a station services more lines, it matters more where a train comes from when determining where it will go next, leading to more dependencies. Indeed, the number of lines station v services is correlated with the size of $\delta(v)$ (0.638 Pearson correlation coefficient).

2) Geographic Node Classification investigates whether EVO captures geographic information. Unlike the London Tube network, where edges connect nearby stations, flights often directly cross geographic regions (e.g., LAX to JFK), so the network does not reflect geography. Itineraries often end in the same region where they began, so memory matters. The task is to predict if an airport is in the Western or Eastern US. **Setup.** We use the Census Bureau’s 9 US divisions [13] and use divisions 1, 2, 3, 5, and 6 (65.7% of airports) as the Eastern region. The task is not structural, so we use `node2vec` in **S2**.

Results. EVO outperforms all baselines by a statistically significant amount, suggesting that EVO captures geographic information in trajectories, which is lost by networks where edges do not respect geography. Indeed, 84.5% of edges cross regions, but 92.8% of itineraries end in the same region as they started (largely due to round-trip flights).

3) Structural Node Classification seeks to predict whether or not an airport is a hub, as characterized by the FAA [14], which defines hubs in terms of the quantity of enplanements. This is highly related to node degree, and is hence *structural*. **Setup.** To capture the structure of nodes, we use `xNetMF` (we also tried `struc2vec`, but saw similar results). The baseline accuracy is 62%, as 38% of airports are not hubs.

Results. SNE performs the best, suggesting that degree information is lost in a HON. Indeed, the median size of $\delta(v)$ for airports is 14. However, for the 25 highest degree airports, the median size is 88. Thus, because hubs have more nuanced flying patterns, they also have more dependencies. The hub’s degree in G is then distributed across more nodes in H (i.e., all those in $\delta(v)$), leading to loss of degree information.

4) Neighborhood Node Classification tests whether dependencies are appropriate for tasks related to node *neighborhoods*. The task is to predict whether a London Tube stop is in an outer or an inner zone. The Tube system is split into 9 zones forming rings around the city center (i.e., zone 1 is the city center, zone 2 is a ring around zone 1, etc.).

Setup. To capture neighborhood information, we use our methods with `node2vec`. Note that 65.1% of stations are in the inner zone (which we consider to be zones 1, 2 and 3).

Results. Either SNE or SWNE perform the best on this task, likely because Tube trajectories cross 3.2 zones on average, biasing node contexts away from neighborhoods. In contrast, a station is generally in the same zone as its neighbors.

E. Clustering

We further investigate RQ1 and RQ2 in an unsupervised setting with more than two classes by clustering nodes and comparing to four ground truth groupings: (1) Tube stations grouped by the line they are on, (2) airports grouped by their region, (3) airports grouped by their FAA hub category, and (4) Tube stations grouped by the zone they are in.

Setup. When using lines as ground truth, we disregard stations which fall on multiple lines, leaving 190 stations. Likewise, when considering zones, we disregard stations in multiple zones, leaving 242. There are 11 Tube lines, 9 zones, 9 airport regions, and 4 airport hub categories (large, medium, small, and non-hub); we use k -Means for clustering with k set to the number of ground truth groups. We use `node2vec` in all tasks except hubs, where we use `struc2vec` (omitting `xNetMF` for brevity). There is no training set, so we try the same `node2vec` (p, q) pairs as classification, and report results with the best pair. We evaluate the quality of clusterings using Normalized Mutual Information (NMI), and report results in Table III.

Results. EVO-max and EVO-avg lead to clusters most similar to station lines and airport regions respectively, while either SNE or SWNE lead to clusters most similar to hubs and

TABLE III: Node Classification (§ IV-D): TASK 1: Classify whether a stop services 1 or >1 line (trajectories matter). Results suggest: EVO captures high-order dependencies. TASK 2: Classify airports by region (geography matters). Results suggest: EVO captures geography via dependencies. TASK 3: Classify airports as hubs (structure matters). Results suggest: degree is lost in H . TASK 4: Classify tube stops as inner/outer-zone (neighborhoods matter). Results suggest: trajectories can be deceptive. **Clustering** (§ IV-E): Results are consistent with classification (except SWNE beats SNE on zones), thus our observations hold even in an *unsupervised* setting with all categories. ^{*} marks results where EVO is stat. sig. better than all baselines or where a baseline is stat. sig. better than EVO-max and EVO-avg.

	Task	Metric	SNE	SWNE	ARWE	EVO-avg	EVO-max
Node Classification	1. Trajectory Node Classification	ACC	0.8259 ± 0.03	0.8481 ± 0.05	0.8506 ± 0.04	0.8321 ± 0.04	0.8901* ± 0.03
		AUC	0.8859 ± 0.04	0.9037 ± 0.04	0.9026 ± 0.04	0.8823 ± 0.05	0.9627* ± 0.02
		F1	0.8792 ± 0.02	0.8960 ± 0.03	0.8969 ± 0.03	0.8886 ± 0.03	0.9254* ± 0.02
	2. Geographic Node Classification	ACC	0.6000 ± 0.06	0.5980 ± 0.07	0.6160 ± 0.04	0.7280* ± 0.05	0.7100* ± 0.05
		AUC	0.5233 ± 0.09	0.5951 ± 0.05	0.5179 ± 0.06	0.7791* ± 0.04	0.7611* ± 0.06
		F1	0.2155 ± 0.09	0.2921 ± 0.12	0.1241 ± 0.13	0.4988* ± 0.09	0.4180* ± 0.08
	3. Structural Node Classification	ACC	0.8748* ± 0.04	N/A	N/A	0.7730 ± 0.05	0.8107 ± 0.04
		AUC	0.9227* ± 0.04	N/A	N/A	0.8785 ± 0.03	0.9006 ± 0.03
		F1	0.8991* ± 0.03	N/A	N/A	0.8212 ± 0.04	0.8429 ± 0.04
	4. Neighborhood Node Classification	ACC	0.9556* ± 0.02	0.9333 ± 0.02	0.9160 ± 0.03	0.9296 ± 0.03	0.9148 ± 0.03
		AUC	0.9916* ± 0.01	0.9867 ± 0.01	0.9791 ± 0.01	0.9770 ± 0.02	0.9684 ± 0.03
		F1	0.9668* ± 0.01	0.9502 ± 0.02	0.9369 ± 0.02	0.9469 ± 0.02	0.9358 ± 0.02
Clustering	1. Lines as ground truth	NMI	0.6478 ± 0.01	0.6258 ± 0.02	0.6458 ± 0.01	0.7100* ± 0.01	0.7125* ± 0.02
	2. Regions as ground truth	NMI	0.1441 ± 0.00	0.1245 ± 0.00	0.1096 ± 0.01	0.3856* ± 0.00	0.2880* ± 0.02
	3. Hubs as ground truth	NMI	0.4138* ± 0.00	N/A	N/A	0.0505 ± 0.01	0.2844 ± 0.00
	4. Zones as ground truth	NMI	0.2542 ± 0.02	0.2705* ± 0.02	0.1933 ± 0.03	0.2192 ± 0.01	0.2444 ± 0.02

station zones. The results are consistent with the classification results, suggesting that our observations hold even in the non-binarized, unsupervised clustering setting.

F. Discussion

For trajectory related tasks, such as clustering stations by line or classifying an airport’s region, we see that EVO provides significant improvement over the baselines, demonstrating that EVO captures high-order dependencies. For instance, high-order dependencies capture information about Tube lines, since where passengers go next depends on what line they are on. Furthermore, while edges in an airport network violate geography, high-order dependencies remember the region of origin, since passengers are likely to return there.

However, for many tasks, accurately capturing high-order dependencies can distract from the information that is actually important for that task. We see this in the degradation of performance on structural tasks, where degree information is lost, and likewise on neighborhood tasks, where node contexts are biased *away* from their neighbors.

We found that either element-wise max or average worked best for f in **S3**. We conjecture that when max works best, features of large magnitude for *any* of the embeddings of $\delta(v)$ are likely also important for v ; max preserves this. When average is best, distributional information may be important. This is reminiscent of max- or average-pooling, e.g. [15].

We conjecture that ARWE’s lower performance is due to the fact that it aggregates context from variable-order random walks *before* learning. In contrast, EVO builds context for each node in V_H separately and learns from that.

V. CONCLUSION

We propose a modular representation learning framework called EVO, which captures non-Markovian dependencies in node embeddings. We investigate EVO’s performance on a wide range of tasks, empirically evaluating when accurately capturing non-Markovian dependencies is useful for network tasks. We find that if a task is related to trajectories, capturing

the non-Markovian dependencies is highly useful, and EVO is an effective way of doing this. However, if a task is unrelated to trajectories, learning over the non-Markovian dependencies can actually be distracting, leading to decreased performance.

Acknowledgements. We thank Leul Wuletaw Belayneh for his early contributions, and Di Jin and Mark Heimann for their helpful discussions. This material is based upon work supported by the National Science Foundation under Grant No. IIS 1845491, an Adobe Digital Experience research faculty award, and an Amazon faculty award.

REFERENCES

- [1] Ryan A Rossi, Nesreen K Ahmed, and Eunye Koh. Higher-order network representation learning. In *WWW (companion)*, 2018.
- [2] Jian Xu, Thanuka L Wickramaratne, and Nitesh V Chawla. Representing higher-order dependencies in networks. *Sci Adv*, 2(5), 2016.
- [3] Ingo Scholtes. When is a network a network?: Multi-order graphical model selection in pathways and temporal networks. In *KDD*, 2017.
- [4] Tara Safavi, Maryam Davoodi, and Danai Koutra. Career transitions and trajectories: A case study in computing. In *KDD*. ACM, 2018.
- [5] Daokun Zhang, Jie Yin, Xingquan Zhu, and Chengqi Zhang. Network representation learning: A survey. *IEEE trans. on Big Data*, 2018.
- [6] Martin Rosvall, Alcides V Esquivel, Andrea Lancichinetti, Jevin D West, and Renaud Lambiotte. Memory in network flows and its effects on spreading dynamics and community detection. *Nature comm*, 5, 2014.
- [7] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, 2013.
- [8] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. Deepwalk: Online learning of social representations. In *KDD*. ACM, 2014.
- [9] Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. In *KDD*. ACM, 2016.
- [10] Leonardo FR Ribeiro, Pedro HP Saverese, and Daniel R Figueiredo. struc2vec: Learning node representations from structural identity. In *KDD*. ACM, 2017.
- [11] Mark Heimann, Haoming Shen, Tara Safavi, and Danai Koutra. Regal: Representation learning-based graph alignment. In *CIKM*. ACM, 2018.
- [12] Jiezhong Qiu, Yuxiao Dong, Hao Ma, Jian Li, Kuansan Wang, and Jie Tang. Network embedding as matrix factorization: Unifying deepwalk, line, pte, and node2vec. In *WSDM*. ACM, 2018.
- [13] Census regions and divisions of the united states. https://www2.census.gov/geo/pdfs/maps-data/maps/reference/us_regdiv.pdf.
- [14] Airport categories. https://www.faa.gov/airports/planning_capacity/passenger_allcargo_stats/categories/, 2018.
- [15] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *NeurIPS*, 2017.