

Higher-Order Ranking and Link Prediction: From Closing Triangles to Closing Higher-Order Motifs

Ryan A. Rossi
Adobe Research
rossi@adobe.com

Anup Rao
Adobe Research
anuprao@adobe.com

Sungchul Kim
Adobe Research
sukim@adobe.com

Eunye Koh
Adobe Research
eunye@adobe.com

Nesreen K. Ahmed
Intel Labs
nesreen.k.ahmed@intel.com

Gang Wu
Adobe Research
gang.wu@adobe.com

ABSTRACT

In this paper, we introduce the notion of *motif closure* and describe higher-order ranking and link prediction methods based on the notion of *closing higher-order network motifs*. The methods are fast and efficient for real-time ranking and link prediction-based applications such as web search, online advertising, and recommendation. In such applications, real-time performance is critical. The proposed methods do not require any explicit training data, nor do they derive an embedding from the graph data, or perform any explicit learning. Most existing methods with the above desired properties are all based on closing triangles (common neighbors, Jaccard similarity, and the ilk). In this work, we investigate higher-order network motifs and develop techniques based on the notion of closing higher-order motifs that move beyond closing simple triangles. All methods described in this work are fast with a runtime that is sublinear in the number of nodes. The experimental results indicate the importance of closing higher-order motifs for ranking and link prediction applications. Finally, the proposed notion of higher-order motif closure can serve as a basis for studying and developing better ranking and link prediction methods.

KEYWORDS

Higher-order motif closure, higher-order ranking, online link prediction, network motifs, graphlets, motif closure, link estimation, online recommendation

1 INTRODUCTION

Link prediction generally refers to predicting the existence of edges (node pairs) in G such that the predicted edges (node pairs) are not in the original edge set E of G . The goal of this task may be to predict future links at time $t + 1$ or to simply predict links that were not observed (*e.g.*, to improve the quality of downstream tasks) [9]. Notice that nearly all link prediction methods first compute a weight $W_{ij} = f(i, j)$ between node i and j and then use W_{ij} to decide whether to predict a link (i, j) or not. We denote the task of estimating a weight $W_{ij} = f(i, j)$ between node i and j as *link weighting* or *link strength estimation*. The weights are then used to derive a ranking of potential links. The potential links may refer to items j that a user i is likely to purchase, or songs that a user is likely to prefer, and so on. In this work, we focus on fast




Figure 1: Higher-Order Motif Closures. The unshaded/white nodes are node i and j . Given a node pair $(i, j) \notin E$ (unshaded/white nodes) and any motif/induced subgraph H , the “edge” between i and j (dotted gray line) is said to close an instance F of H if the edge (i, j) were to actually exist in G .

and efficient methods for computing link weights based on closing higher-order network motifs. Such weights based on higher-order motif closures can then be used for ranking-based applications (such as recommender systems and the ilk).

Ranking is a key component of many real-world applications such as web search, online advertising, and recommendation [7]. In these applications, real-time performance is critical, *e.g.*, in web search users expect an answer to their query in the order of a few hundred milliseconds [2, 4]. This makes it impossible to learn a complex ranking function. Instead, there are usually two components to such a system. In the first component, a *fast online approach* is used to identify the top- k most relevant results in real-time (where k is typically small), which are then displayed to the user. In the second component, a more accurate but computationally expensive model is trained to improve the initial ranking. The ranking learned from the model can be used directly or combined with simpler approaches to obtain a final re-ranking of the web pages (or items). In this work, we primarily focus on the first component.

Ranking and link prediction [4] are important fundamental problems with many applications including recommendation of items [4], friends [6], web pages [7], among many others [7, 9]. Common neighbors and approaches based on common neighbors such as Jaccard similarity are known to be strong baselines that are hard to beat in practice [11]. These baselines are all fundamentally based on the notion of “closing triangles” [1, 9]. They are both simple and fast for ranking in an online real-time fashion. In this work, we investigate higher-order network motifs and develop ranking techniques based on the notion of *closing higher-order motifs* (Definition 1) that move beyond “closing” simple triangles.

Table 1: Mean average precision (MAP) results for ranking (and prediction) methods based on closing higher-order motifs.

	<i>bn-mouse</i>	<i>bio-DM-LC</i>	<i>bio-CE-HT</i>	<i>bio-DM-HT</i>	<i>ia-reality</i>	<i>web-polblogs</i>	<i>biogrid-worm</i>	<i>biogrid-plant</i>	<i>biogrid-yeast</i>	<i>email-dnc-corec.</i>	<i>soc-advogato</i>	<i>econ-wrml</i>	<i>bn-macaque-the.</i>	<i>road-minnesota</i>	<i>soc-fb-messages</i>	<i>email-EU</i>	<i>email-univ</i>
4-path	0.829	0.687	0.607	0.594	0.649	0.778	0.865	0.729	0.893	0.873	0.914	0.788	0.942	0.326	0.844	0.854	0.707
4-star	0.880	0.787	0.595	0.696	0.922	0.814	0.895	0.861	0.840	0.813	0.889	0.688	0.961	0.388	0.807	0.972	0.695
4-cycle	0.881	0.958	0.651	0.926	0.827	0.885	0.908	0.935	0.927	0.957	0.930	0.900	0.773	0.950	0.870	0.902	0.847
4-tailed-triangle	0.804	0.612	0.570	0.752	0.773	0.663	0.773	0.681	0.689	0.779	0.600	0.496	0.530	0.834	0.722	0.937	0.582
4-chordal-cycle	0.801	0.837	0.598	0.842	0.312	0.966	0.840	0.854	0.977	0.996	0.986	0.947	0.750	0.939	0.935	0.782	0.969
4-clique	0.804	0.838	0.595	0.843	0.293	0.963	0.842	0.847	0.972	0.997	0.986	0.965	0.759	0.939	0.960	0.798	0.982
CN	0.705	0.872	0.613	0.839	0.422	0.814	0.833	0.897	0.839	0.960	0.949	0.852	0.342	0.945	0.790	0.890	0.941
Jaccard Sim.	0.705	0.873	0.618	0.841	0.537	0.933	0.853	0.918	0.955	0.997	0.973	0.918	0.764	0.944	0.841	0.933	0.949
Adamic/Adar	0.705	0.883	0.621	0.842	0.549	0.940	0.856	0.920	0.959	0.997	0.976	0.919	0.777	0.945	0.848	0.935	0.953

While most existing work focuses on learning a ranking function [3, 5, 12], we instead focus on direct principled approaches that are efficient (sublinear in the number of nodes), can be directly computed in real-time, easily parallelizable, and naturally amenable for online real-time ranking in the streaming setting. This work introduces the general notion of closing higher-order motifs and based on this notion we develop direct ranking techniques that are efficient for real-time online ranking and prediction. Compared to similar techniques that can be used for this setting such as Common Neighbors and methods based on it (e.g., Jaccard similarity), the proposed techniques are fundamentally more powerful as they naturally generalize over these existing techniques that are all based on closing triangles (a lower-order motif). The proposed notion of higher-order motif closure can serve as a basis for studying and developing better ranking (and prediction) methods based on the higher-order motif closures.

2 CLOSING HIGHER-ORDER MOTIFS

We first introduce the notion of a higher-order *network motif closure* that lies at the heart of this work.

DEFINITION 1 (MOTIF CLOSURE). A node pair (i, j) is said to close a network motif H iff adding an edge (i, j) to E closes an instance $F \in I'_G(H)$ of motif H where $G' = (V, E \cup \{(i, j)\})$ and $I'_G(H)$ is the set of unique instances of motif H in G' .

Figure 1 provides a few examples of higher-order motif closures. The edge (i, j) shown as a dotted line in Figure 1 closes each motif. For instance, the edge between node i and j in the rightmost motif in Figure 1 closes a 4-clique. We now formally introduce the frequency of higher-order motif closures for a node pair (i, j) as follows:

DEFINITION 2 (HIGHER-ORDER MOTIF CLOSURE FREQUENCY). Let $G' = (V, E')$ where $E' = E \cup \{(i, j)\}$ and let $I'_G(H)$ be the set of unique instances of motif H in G' . Then the frequency of closing a higher-order motif H between node i and j is:

$$W_{ij} = \sum_{F \in I'_{G'}(H)} \mathbb{I}(\{(i, j)\} \in E'(F)) \quad (1)$$

where W_{ij} is equal to the number of unique instances of H that contain nodes $\{i, j\} \subset V(G')$ as an edge.

We provide a simple routine in Algorithm 1 for computing the weight W_{ij} representing the frequency of closing motif H between node i and j . The approach has two simple steps. First, given an arbitrary node pair (i, j) , a motif H of interest, and the current graph $G = (V, E)$, we simply add the node pair (i, j) as an edge by setting $E' \leftarrow E \cup \{(i, j)\}$ and $G' = (V, E')$ (Alg. 1 Line 1).¹ As an aside, this can be performed *implicitly* without any additional work. However, it is shown in Algorithm 1 since after adding (i, j) to the edge set, we can use the fastest known algorithm for counting the occurrences of motif (induced subgraph/graphlet) H between node i and j in G' . Nevertheless, we can always modify the best known algorithm [1] so that it *implicitly* treats the pair of nodes (i, j) given as input as being connected for the sake of determining the number of instances of H that would be closed if (i, j) were to really exist as an edge in G . Second, we compute the number of instances of motif H that contain nodes i and j in G' (Alg. 1 Line 2). Given a set $\mathcal{Y} = \{y_1, y_2, \dots, y_j, \dots\}$ of nodes (items, ads, songs, friends) to be ranked, Algorithm 1 can be used to obtain $W_{ij} = f(x_i, y_j)$, $\forall j = 1, \dots, |\mathcal{Y}|$.

Algorithm 1 Higher-Order Motif Closures

Input: a graph $G = (V, E)$, node pair (i, j) , and network motif/graphlet H
Output: the frequency W_{ij} of motif closures of H for nodes i and j

- 1 Set $E' \leftarrow E \cup \{(i, j)\}$ and $G' = (V, E')$
- 2 Use fast algorithm [1, 10] to compute $W_{ij} = \#$ of occurrences of motif H between node i and j in G'

Extending Other Measures using Motif Closure. Given two nodes i and j , Common Neighbor-based methods are those that use the quantity $|\Gamma_i \cap \Gamma_j|$ where Γ_i and Γ_j are the set of neighbors for node i and j , respectively. Common neighbors is simply $W_{ij} = |\Gamma_i \cap \Gamma_j|$

¹Note that if edges are arriving continuously over time in a streaming fashion, then we may also encounter a node i (or j) such that $i \notin V$. In this case, we also set $V' \leftarrow V \cup \{i\}$ and $G' = (V', E')$.

Table 2: Coverage (\downarrow) results for the ranking methods based on closing higher-order network motifs. Lower is better.

	<i>bn-mouse</i>	<i>bio-DM-LC</i>	<i>bio-CE-HT</i>	<i>bio-DM-HT</i>	<i>ia-reality</i>	<i>web-polblogs</i>	<i>biogrid-worm</i>	<i>biogrid-plant</i>	<i>biogrid-yeast</i>	<i>email-dnc-corec.</i>	<i>soc-advogato</i>	<i>econ-wm1</i>	<i>bn-macaque-rhe.</i>	<i>road-minnesota</i>	<i>soc-fb-messages</i>	<i>email-EU</i>	<i>email-univ</i>
4-path	0.606	0.964	0.822	0.962	0.537	0.958	0.828	0.980	0.963	0.911	0.936	0.976	1	0.999	0.970	0.579	0.991
4-star	0.637	0.950	0.815	0.944	0.024	0.911	0.972	0.963	0.945	0.998	0.985	0.915	0.088	0.999	0.953	0.245	0.986
4-cycle	0.300	0.375	0.922	0.645	0.457	0.942	0.542	0.483	0.869	0.801	0.952	0.942	0.995	1	0.971	0.366	0.897
4-tailed-triangle	1	0.910	1	1	0.451	0.967	0.654	0.756	0.986	0.893	0.997	0.964	1	1	0.981	0.404	0.992
4-chordal-cycle	1	0.913	1	0.913	0.936	0.613	1	1	0.228	0.211	0.217	0.796	0.965	1	1	0.620	0.390
4-clique	1	1	1	1	1	0.709	1	1	0.322	0.216	0.315	0.521	0.989	1	0.902	0.62	0.409
CN	1	0.732	1	0.826	0.652	0.705	0.752	0.739	0.788	0.224	0.782	0.894	0.972	1	0.864	0.461	0.766
Jaccard Sim.	1	0.732	1	0.826	0.658	0.707	0.752	0.739	0.824	0.225	0.782	0.902	0.960	1	0.885	0.462	0.762
Adamic/Adar	1	0.732	1	0.826	0.653	0.706	0.747	0.739	0.823	0.218	0.783	0.905	0.923	1	0.866	0.462	0.761




Figure 2: Precision at $k = 1, \dots, 40$ for different network motif closure rankings.

where W_{ij} represents the number of *potential triangles* that would be closed if there were an edge between i and j . The notion of “closing” triangles lies at the heart of many other existing methods that are based on $|\Gamma_i \cap \Gamma_j|$ such as Jaccard similarity, Adamic/Adar (AA), among others. All of these methods can be viewed as extensions of Common Neighbors with some form of normalization, e.g., Jaccard similarity is $W_{ij} = |\Gamma_i \cap \Gamma_j| / |\Gamma_i \cup \Gamma_j|$. Extending the proposed higher-order motif-based link ranking and prediction techniques is left for future work. This includes extending the notion of “closing” higher-order network motifs for other measures such as Jaccard similarity, Adamic/Adar, among any others where the notion of *closing triangles* can be replaced with the notion of *closing a higher-order motif* introduced in this work.

3 EXPERIMENTS

The experiments are designed to evaluate the effectiveness of the proposed methods that are based on the notion of “closing” higher-order network motifs. These methods go beyond closing simple triangles. To ensure the significance and generality of our findings (as much as possible), we evaluate the proposed methods using a wide variety of networks from different application domains. All data was obtained from NetworkRepository [8].

We compare the proposed higher-order motif closure methods against CN-based methods (CN, Jaccard similarity, Adamic/Adar) since these are all based on closing triangles also have the same

desired properties as the higher-order motif closure methods described in this paper. In this work, we only investigate the most basic and fundamental higher-order motif closures. Developing more sophisticated higher-order ranking measures based on these fundamental motif closures is left for future work. However, we did run a few experiments using an extended higher-order Jaccard similarity (one for each motif closure, giving 6 total for 4-node motifs) and higher-order Adamic/Adar ranking measures, again giving 6 new rankings total. Since each variant provides 6 additional rankings, the results were removed for brevity, but in some cases performed better than the most basic motif closures introduced in this paper. As such, the proposed notion of higher-order motif closures serve as fundamental building blocks for developing better higher-order ranking and prediction methods.

Unless otherwise mentioned, we hold-out 10% of the observed node pairs and randomly sample the same number of negative node pairs. We then use the methods to obtain a ranking of the node pairs in this set.² Recall the proposed techniques do not require learning a sophisticated model nor do they require training data. As such, the notion of motif closure proposed in this work can be used in a real-time streaming fashion and has many obvious advantages to more sophisticated model-based approaches. Mean

²In recommender systems, the set of node pairs to be ranked is actually a smaller set of “relevant items” $\mathcal{Y}_i = \{y_1, \dots, y_j, \dots\} \subset \mathcal{Y}$ for a user i . Nevertheless, this can also be viewed as a ranking of node pairs where user i is fixed.

Table 3: Robustness results (MAP). See text for discussion.

	<i>bn-mouse</i>	<i>bio-DM-LC</i>	<i>bio-CE-HT</i>	<i>bio-DM-HT</i>	<i>ia-reality</i>	<i>web-pollblogs</i>	<i>biogrid-worm</i>	<i>biogrid-plant</i>
4-path	0.764	0.698	0.495	0.582	0.774	0.812	0.899	0.847
4-star	0.868	0.781	0.52	0.645	0.920	0.81	0.894	0.836
4-cycle	0.763	0.915	0.482	0.871	0.665	0.855	0.863	0.872
4-tailed-triangle	0.684	0.700	0.444	0.755	0.725	0.708	0.765	0.658
4-chordal-cycle	0.781	0.816	0.440	0.811	0.237	0.942	0.812	0.776
4-clique	0.781	0.825	0.440	0.809	0.205	0.944	0.807	0.761
CN	0.686	0.800	0.458	0.801	0.283	0.806	0.791	0.819
Jaccard Sim.	0.693	0.809	0.461	0.803	0.357	0.910	0.808	0.828
Adamic/Adar	0.704	0.809	0.462	0.803	0.375	0.912	0.808	0.831

Average Precision (MAP) results are provided in Table 1 whereas coverage is provided in Table 2.

CONCLUSION 1. Ranking based on closing higher-order motifs outperforms other direct methods that are based on closing triangles.

In nearly all cases, the higher-order motif closures achieve better precision and coverage than techniques based on closing lower-order triangles.

CONCLUSION 2. The best performing motif closure is consistent across different evaluation measures. The motif closure that achieves the best precision (Table 1) is typically the same motif that achieves the best coverage (Table 2).

CONCLUSION 3. There is no single higher-order motif closure that performs best for all graphs. The best motif depends highly on the structural characteristics of the graph and its domain (biological vs. social network) as shown in Table 1 and Table 2.

In Table 1-2, biological and brain networks achieve best performance using the ranking given by 4-cycle and 4-star closures. This also holds true for the interaction (ia-reality) and road network investigated. The 4-star and 4-cycle motif closures are more sparse compared to the 4-chordal-cycle (paw motif) and 4-clique motif closure. In the web graph, economic, and social networks, both the 4-chordal-cycle (diamond motif closure) and 4-clique motif closure achieves significantly better performance than the other motif closures. Notice that both these motif closures are composed of two or more triangles and thus can be seen as a stronger triadic closure motif. The 4-path, 4-tailed-triangle, and triangle (CN) motif closures did not perform the best in any of the graphs investigated. That is, there were always a higher-order motif closure with better performance as shown in Table 1 and Table 2. In Figure 2, we also show the precision at $k = 1, \dots, 40$ for closing different higher-order network motifs. In nearly all cases, the rankings given by the 4-node motif closures are better than the lower-order CN approach that is based on closing triangles.

Robustness of Ranking from Higher-Order Motif Closures. In addition, we investigate the robustness of the higher-order motif closures to noise in the graph, *i.e.*, random link additions. To understand the robustness of the motif closure methods for graphs

with noisy and spurious links, we select pairs of nodes uniformly at random that are not linked in G and create a link between each pair. In this set of experiments, we sample $|E|/2$ node pairs (negative/unobserved edges) and add them to G . Results are shown in Table 3. Due to space, we show only a subset of the networks used in Table 1-2.

CONCLUSION 4. Robustness of the ranking by higher-order motif closures is comparable, and slightly better than techniques based on closing triangles.




Figure 3: Average runtime in milliseconds to compute all motif closures for each node pair.

Runtime performance. We report the average runtime in milliseconds to compute all motif closures for each node pair in G . For most graphs, it takes less than a millisecond on average as shown in Figure 3 and therefore is fast for large-scale ranking problems.

CONCLUSION 5. For any 4-node motif H , counting the number of motif closures W_{ij} that would arise if (i, j) was added to G is fast taking less than a millisecond on average (across all graphs).

The runtime can be significantly improved for certain problem settings: Suppose we are interested in only the top- k most relevant node pairs (or items for a user i) given by a ranking from an arbitrary motif closure for motif H , then for possibly many such node pairs, we can avoid computing W_{ij} (*i.e.*, # of instances of motif H in G that would be closed if the node pair (i, j) actually existed/observed in G) altogether by first deriving an upper bound UB of W_{ij} in $o(1)$ constant time and only computing W_{ij} if $UB > \delta$ where δ is the weight of the node pair in the top- k ranking with minimum weight (the node pair with rank k). Since otherwise we know W_{ij} is not large enough to beat the node pair with the k -th largest weight.

4 CONCLUSION

This work proposed the notion of *motif closure* and described higher-order ranking and link prediction techniques based on the notion of *closing higher-order network motifs*. Such techniques were shown to be effective for online real-time ranking (and prediction) as they often outperformed a number of baselines that are based on closing triangles. Future work will investigate using the notion of *closing higher-order motifs* to extend other techniques such as a higher-order Jaccard similarity or higher-order Adamic/Adar measures

From Closing Triangles to Closing Higher-Order Motifs

based on closing higher-order network motifs such as 4-cliques, 4-cycles, among others.

REFERENCES

- [1] Nesreen K. Ahmed, Jennifer Neville, Ryan A. Rossi, and Nick Duffield. 2015. Efficient Graphlet Counting for Large Networks. In *ICDM*. 10.
- [2] Michael W Berry and Murray Browne. 2005. *Understanding Search Engines: Mathematical Modeling and Text Retrieval*. Vol. 17. SIAM.
- [3] Sougata Chaudhuri and Ambuj Tewari. 2015. Online ranking with top-1 feedback. In *Artificial Intelligence and Statistics*. 129–137.
- [4] Mukund Deshpande and George Karypis. 2004. Item-based top-n recommendation algorithms. *TOIS* 22, 1 (2004), 143–177.
- [5] Ernesto Diaz-Aviles, Lucas Drumond, Lars Schmidt-Thieme, and Wolfgang Nejdl. 2012. Real-time top-n recommendation in social streams. In *RecSys*. ACM, 59–66.
- [6] David Liben-Nowell and Jon Kleinberg. 2007. The link-prediction problem for social networks. *JASIST* 58, 7 (2007), 1019–1031.
- [7] Christopher Manning, Prabhakar Raghavan, and Hinrich Schütze. 2010. Introduction to information retrieval. *Nat. Lang. Eng.* 16, 1 (2010), 100–103.
- [8] Ryan A. Rossi and Nesreen K. Ahmed. 2015. The Network Data Repository with Interactive Graph Analytics and Visualization. In *AAAI*. 4292–4293. <http://networkrepository.com>
- [9] Ryan A. Rossi, Luke K. McDowell, David W. Aha, and Jennifer Neville. 2012. Transforming graph data for statistical relational learning. *JAIR* (2012), 363–441.
- [10] Ryan A. Rossi, Rong Zhou, and Nesreen K. Ahmed. 2018. Estimation of Graphlet Counts in Massive Networks. In *TNNLS*. 44–57.
- [11] Anton Tsitsulin, Davide Mottin, Panagiotis Karras, and Emmanuel Müller. 2018. Verse: Versatile graph embeddings from similarity measures. In *WWW*. 539–548.
- [12] Hyokun Yun, Parameswaran Raman, and S Vishwanathan. 2014. Ranking via robust binary classification. In *NIPS*. 2582–2590.