
Motif Prediction with Graph Neural Networks

Maciej Besta¹ Raphael Grob¹ Cesare Miglioli² Nicola Bernold¹ Grzegorz Kwasniewski¹

Gabriel Gjini¹ Raghavendra Kanakagiri³ Saleh Ashkboos¹ Lukas Gianinazzi¹

Nikoli Dryden¹ Torsten Hoefler¹

¹Department of Computer Science, ETH Zurich

²Research Center for Statistics, University of Geneva

³Department of Computer Science, Indian Institute of Technology Tirupati

Abstract

Link prediction is one of the central problems in graph mining. However, recent studies highlight the importance of *higher-order network analysis*, where complex structures called motifs are the first-class citizens. We first show that existing link prediction schemes fail to effectively predict motifs. To alleviate this, we establish a general *motif prediction problem* and we propose several heuristics that assess the chances for a specified motif to appear. To make the scores realistic, our heuristics consider – among others – *correlations between links*, i.e., the potential impact of some arriving links on the appearance of other links in a given motif. Finally, for highest accuracy, we develop a graph neural network (GNN) architecture for motif prediction. Our architecture offers vertex features and sampling schemes that capture the rich structural properties of motifs. While our heuristics are fast and do not need any training, GNNs ensure highest accuracy of predicting motifs, both for dense (e.g., k -cliques) and for sparse ones (e.g., k -stars). We consistently outperform the best available competitor by more than 10% on average and up to 32% in area under the curve. Importantly, the advantages of our approach over schemes based on uncorrelated link prediction increase with the increasing motif size and complexity. We also successfully apply our architecture for predicting more arbitrary *clusters* and *communities*, illustrating its potential for graph mining beyond motif analysis.

1 Introduction and Motivation

One of the central problems in graph mining and learning is link prediction [5, 6, 59, 76, 93, 95], in which one is interested in predicting if a given *pair of vertices* will become connected. However, recent works argue the importance of *higher-order graph organization* [10], where one focuses on finding and analyzing small recurring *subgraphs* called *motifs* (sometimes referred to as *graphlets* or *graph patterns*) instead of individual links. Motifs are central to many graph mining problems in computational biology, chemistry, and a plethora of other fields [16, 22, 23, 28, 30, 41, 46]. Specifically, motifs are building blocks of different networks, including transcriptional regulation graphs, social networks, brain graphs, or air traffic patterns [10]. There exist many motifs, for example k -cliques, k -stars, k -clique-stars, k -cores, and others [20, 43, 54]. A huge number of works are dedicated to motif *counting*, *listing* (also called *enumeration*), or *checking for the existence* of a given motif [22, 30]. However, while a few recent schemes focus on predicting *triangles* [9, 63, 64], no works target the problem of *general motif prediction*, i.e., analyzing whether specified complex structures may appear in the data. As with link prediction, it would enable predicting the evolution

of data, but also finding missing structures in the available data. For example, one could use motif prediction to find probable missing clusters of interactions in biological (e.g., protein) networks, and use the outcomes to limit the number of expensive experiments conducted to find missing connections [59, 60].

In this paper, we first (Section 3) establish and formally describe a general motif prediction problem, going beyond link prediction and illustrating how to predict future (or missing from the data) higher-order network patterns. A key challenge is the appropriate *problem formulation*. Similarly to link prediction, one wants a *score function* that – for a given vertex set V_M – assesses the chances for a given motif to appear. However, the function must consider the increased complexity of the problem (compared to link prediction). In general, contrary to a single link, a motif may be formed by an *arbitrary* number $|V_M|$ of vertices, and the number of potential edges between these vertices can be large, i.e., $O(|V_M|^2)$. For example, one may be interested in analyzing whether a *group* of entities V_M may become a k -clique in the future, or whether a specific vertex $v \in V_M$ will become a *hub* of a k -star, connecting v to $k - 1$ other selected vertices from $V_M \setminus \{v\}$. This leads to novel issues, not present in link prediction. For example, what if *some* edges, belonging to the motif being predicted, already exist? How should they be treated by a score function? Or, how to enable users to apply their domain knowledge? For example, when predicting whether the given vertices will form some chemical particle, a user may know that the presence of some link (e.g., some specific atomic bond) may increase (or decrease) the chances for forming another bond. Now, how could this knowledge be provided in the motif score function? We formally specify these and other aspects of the problem in a general theoretical framework, and we provide example motif score functions. We explicitly consider correlations between edges forming a motif, i.e., the fact that the appearance of some edges may increase or decrease the overall chances of a given motif to appear.

Then, we develop a learning architecture based on graph neural networks (GNNs) to further enhance motif prediction accuracy (Section 4). For this, we extend the state-of-the-art SEAL link prediction framework [93] to support arbitrary motifs. For a given motif M in question, we train our architecture on what is the “right motif surroundings” (i.e., nearby vertices and edges) that could result in the appearance of M . Then, for a given set of vertices V_M , the architecture infers the chances for M to appear. The key

challenge is to be able to capture the *richness* of different motifs and their surroundings. We tackle this with an appropriate selection of *negative* samples, i.e., subgraphs that resemble the searched motifs but that are not identical to them. Moreover, when selecting the size of the “motif surroundings” we rely on an assumption also used in link prediction, which states that only the “close surroundings” (i.e., nearby vertices and edges, 1–2 hops away) of a link to be predicted have a significant impact on whether or not this link would appear [93, 95]. We use this assumption for motifs: as our evaluation shows, it ensures high accuracy while significantly reducing runtimes of training and inference (as only a small subgraph is used, instead of the whole input graph). We call our GNN architecture **SEAM**: learning from Subgraphs, Embeddings and Attributes for **Motif** prediction¹. Our evaluation (Section 5) illustrates the high accuracy of SEAM (often $>90\%$), for a variety of graph datasets and motif sizes.

To motivate our work, we now compare SEAM and a proposed Jaccard-based heuristic that considers link correlations to two baselines that straightforwardly *use link prediction independently for each motif link*: a Jaccard-based score and the state-of-the-art SEAL scheme based on GNNs [93]. We show the results in Figure 1. The correlated Jaccard outperforms a simple Jaccard, while the proposed SEAM is better than SEAL. The benefits generalize to different graph datasets. Importantly, we observe that the larger the motif to predict becomes (larger k), *the more advantages our architecture delivers*. This is because larger motifs provide more room for *correlations between their associated edges*. Straightforward link prediction based schemes do not consider this effect, while our methods

SEAL, Jaccard:		Accuracy decreases with motif size			Accuracy decreases with motif size		
Jaccard		66.39	61.13	59.64	74.31	59.99	53.26
Correlated Jaccard		67.63	64.48	62.52	72.85	64.90	61.19
SEAL		77.61	73.32	72.18	77.50	70.62	65.46
SEAM		86.92	89.48	91.80	90.97	96.33	98.15
		3-star	5-star	7-star	3-clique	5-clique	7-clique
		SEAM: Accuracy increases with motif size			SEAM: Accuracy increases with motif size		

Figure 1: **Motivating our work (SEAM)**: the accuracy (%) of predicting different motifs with SEAM compared to using a state-of-the-art SEAL link prediction scheme [93, 95] and a naive one that does not consider correlations between edges. The details of the experimental setup are in Section 5 (the dataset is USAir). **Importantly**: (1) SEAM outperforms all other methods, (2) the accuracy of SEAM *increases* with the size (k) of each motif, while in other methods it *decreases*.

¹In analogy to SEAL [93, 95], which stands for “learning from Subgraphs, Embeddings, and Attributes for Link prediction”.

do, which is why we offer more benefits for more complex motifs. The advantages of SEAM over the correlated Jaccard show that GNNs more robustly capture correlations and the structural richness of motifs than simple manual heuristics. Simultaneously, heuristics do not need any training.

Finally, SEAM also successfully predicts more arbitrary *communities* or *clusters* [16, 22, 39, 54]. They differ from motifs as they do not have a very specific fixed structure (such as a star) but simply have the edge density above a certain threshold. SEAM’s high accuracy in predicting such structures illustrates its potential for broader graph mining beyond motif analysis.

2 Background and Notation

We first describe the necessary background and notation.

Graph Model We model an undirected graph G as a tuple (V, E) ; V and $E \subseteq V \times V$ are sets of nodes (vertices) and links (edges); $|V| = n$, $|E| = m$. Vertices are modeled with integers $1, \dots, n$; $V = \{1, \dots, n\}$. N_v denotes the neighbors of $v \in V$; $d(v)$ denotes the degree of v .

Link Prediction We generalize the well-known link prediction problem. Consider two unconnected vertices u and v . We assign a *similarity score* $s_{u,v}$ to them. All pairs of vertices that are not edges receive such a score and are ranked according to it. The higher a similarity score is, the “more likely” a given edge is to be missing in the data or to be created in the future. We stress that the link prediction scores are usually not based on any probabilistic notion (in the formal sense) and are only used to make comparisons between pairs of vertices in the same input graph dataset.

There are numerous known similarity scores. First, a large number of scores are called *first order* because they only consider the neighbors of u and v when computing $s_{u,v}$. Examples are the

Common Neighbors scheme $s_{u,v}^{CN} = |N_u \cap N_v|$ or the **Jaccard** scheme $s_{u,v}^J = \frac{|N_u \cap N_v|}{|N_u \cup N_v|}$ [17].

These schemes assume that two vertices are more likely to be linked if they have many common neighbors. There also exist **higher-order** similarity schemes. Such schemes also consider vertices not directly attached to u and v . All these schemes can be described using the same formalism of the γ -*decaying heuristic* proposed by [93]. Intuitively, for a given pair of vertices (u, v) , the γ -decaying heuristic for (u, v) provides a sum of contributions into the link prediction score for (u, v) from all other vertices, weighted in such a way that nearby vertices have more impact on the score.

Graph Neural Networks (GNNs) Graph neural networks (GNNs) are a recent class of neural networks for learning over irregular data such as graphs [26, 29, 72, 73, 78, 88, 90, 91, 96, 97]. There exists a plethora of models and methods for GNNs; most of them consist of two fundamental parts: (1) an aggregation layer that combines the features of the neighbors of each node, for all the nodes in the input graph, and (2) combining the scores into a new score. The input to a GNN is a tuple $G = (A, X)$. The input graph G having n vertices is modeled with an adjacency matrix $A \in \mathbb{R}^{n \times n}$. The features of vertices (with dimension d) are modeled with a matrix $X \in \mathbb{R}^{n \times d}$.

3 Motif Prediction: Formal Statement and Score Functions

We now formally establish the motif prediction problem. We define a motif as a pair $M = (V_M, E_M)$. V_M is the set of *existing* vertices of G that form a given motif ($V_M \subseteq V$). E_M is the set of edges of G that form the motif being predicted; some of these edges may already exist ($E_M \subset V_M \times V_M$).

We make the problem formulation (in § 3.1–§ 3.3) *general*: it can be applied to any graph generation process. Using this formulation, one can then devise specific heuristics that may assume some details on how the links are created, similarly as is done in link prediction. Here, we propose example motif prediction heuristics that harness the Jaccard, Common Neighbors, and Adamic-Adar link scores.

3.1 Motif Prediction vs. Link Prediction

We illustrate the motif prediction problem by discussing the differences between link and motif prediction. Figure 2 provides an overview of our work. We consider all these differences when proposing specific schemes for predicting motifs.

(M) There May Be Many Potential New Motifs for a Fixed Vertex Set Link prediction is a “binary” problem: for a given pair of vertices that are not connected, there can be only one link appearing. In

motif prediction, the situation is more complex. Consider vertices v_1, \dots, v_k . There are many possible motifs to appear between v_1, \dots, v_k . We now state a precise count; the proof is in the appendix.

Observation 1. Consider vertices $v_1, \dots, v_k \in V$. Assuming no edges already connecting v_1, \dots, v_k , there are $2^{\binom{k}{2}} - 1$ motifs (with between 1 and $\binom{k}{2}$ edges) that can appear to connect v_1, \dots, v_k .

(E) Incoming Motifs May Have Existing Edges A link can only appear between *unconnected* vertices. Contrarily, a motif can appear and connect vertices *already* with some edges between them.

(D) There May Be “Deal-Breaker” Edges There may be some edges, the appearance of which would make the appearance of a given motif *unlikely* or even *impossible* (e.g., existing chemical bonds could prevent other bonds). For example, consider a prediction query where one is interested whether a given vertex set can become connected with a *star* but in such a way that *none of the non-central vertices are connected to one another*. Now, if there is already some edge connecting these non-central vertices, this makes it impossible a given motif to appear while satisfying the query. We will refer to such edges as the “deal-breaker” edges.

(L) Motif Prediction Query May Depend on Vertex Labeling The query can depend on a specific vertex labeling. For example, when asking whether a 5-star will connect six given vertices v_1, \dots, v_6 , one may be interested in *any* 5-star connecting v_1, \dots, v_6 , or a 5-star connecting these vertices in a *specific way*, e.g., with its center being v_1 . We enable the user to specify how edges in E_M should connect vertices in V_M .

3.2 Types of Edges in Motifs

We first describe different forms of edges that are related to a motif. First, note that $E_M = E_{M,\mathcal{N}} \cup E_{M,\mathcal{E}}$ where $E_{M,\mathcal{N}}$ are edges that do not exist in G at the moment of querying ($\forall e \in E_{M,\mathcal{N}} e \notin E$; \mathcal{N} indicates “*N*on-existing”) and $E_{M,\mathcal{E}}$ are edges that already exist, cf. (E) in § 3.1 ($\forall e \in E_{M,\mathcal{E}} e \in E$; \mathcal{E} indicates “*E*xisting”). There may also be edges between vertices in V_M which do *not* belong to the motif M (i.e., they belong to $E_{V_M} = \{\{i, j\} : i, j \in V_M \wedge i \neq j\}$ but *not* E_M). We refer to such edges as \overline{E}_M since $E_{V_M} = \overline{E}_M \cup E_M$ (i.e., a union of disjoint sets by definition). Some edges in \overline{E}_M may be deal-breakers (cf. (D) in § 3.1), we denote them as $\overline{E}_{M,\mathcal{D}}$ (\mathcal{D} indicates “*D*eal-breaker”).

Non deal-breakers that are in \overline{E}_M are denoted with $\overline{E}_{M,\mathcal{I}}$ (\mathcal{I} indicates “*I*ntert”). Note that $\overline{E}_M = \overline{E}_{M,\mathcal{D}} \cup \overline{E}_{M,\mathcal{I}}$ and $\overline{E}_M = E_{V_M} \setminus E_M$. To conclude, as previously done for the set E_M , we note that $\overline{E}_{M,\mathcal{D}} = \overline{E}_{M,\mathcal{D},\mathcal{N}} \cup \overline{E}_{M,\mathcal{D},\mathcal{E}}$ where $\overline{E}_{M,\mathcal{D},\mathcal{N}}$ are deal-breaker edges that do not exist in G at the moment of querying ($\forall e \in \overline{E}_{M,\mathcal{D},\mathcal{N}} e \notin E$; \mathcal{N} indicates “*N*on-existing”)

and $\overline{E}_{M,\mathcal{D},\mathcal{E}}$ are deal-breaker edges that already exist, cf. (E) in § 3.1 ($\forall e \in \overline{E}_{M,\mathcal{D},\mathcal{E}} e \in E$; \mathcal{E} indicates “*E*xisting”). We explicitly consider $\overline{E}_{M,\mathcal{D},\mathcal{N}}$ because – even if a given deal-breaker edge does not exist, but it *does* have a large chance of appearing – the overall motif score should become lower.

Symbol	Description
E_M	All edges forming a motif in question; $E_M = E_{M,\mathcal{N}} \cup E_{M,\mathcal{E}}$
$E_{M,\mathcal{N}}$	Motif edges that do not yet exist
$E_{M,\mathcal{E}}$	Motif edges that already exist in the data
\overline{E}_M	Edges not in E_M , defined over vertex pairs in V_M ; $\overline{E}_M = \overline{E}_{M,\mathcal{D}} \cup \overline{E}_{M,\mathcal{I}}$
E_{V_M}	All possible edges between motif vertices; $E_{V_M} = \overline{E}_M \cup E_M$
$\overline{E}_{M,\mathcal{D}}$	Deal-breaker edges; $\overline{E}_{M,\mathcal{D}} = \overline{E}_{M,\mathcal{D},\mathcal{N}} \cup \overline{E}_{M,\mathcal{D},\mathcal{E}}$
$\overline{E}_{M,\mathcal{D},\mathcal{N}}$	Deal-breaker edges that do not exist yet
$\overline{E}_{M,\mathcal{D},\mathcal{E}}$	Deal-breaker edges that already exist
$\overline{E}_{M,\mathcal{I}}$	Non deal-breaker edges in \overline{E}_M ; “edges that do not matter”
E_M^*	“Edges that matter for the score”: $E_M^* = E_M \cup \overline{E}_{M,\mathcal{D}}$
$E_{M,\mathcal{E}}^*$	All existing edges “that matter”: $E_{M,\mathcal{E}}^* = E_{M,\mathcal{E}} \cup \overline{E}_{M,\mathcal{D},\mathcal{E}}$
$E_{M,\mathcal{N}}^*$	All non-existing edges “that matter”: $E_{M,\mathcal{N}}^* = E_{M,\mathcal{N}} \cup \overline{E}_{M,\mathcal{D},\mathcal{N}}$

Table 1: Different types of edges used in this work.

3.3 General Problem and Score Formulation

We now formulate a general *motif prediction score*. Analogously to link prediction, we assign scores to motifs, to be able to quantitatively assess which motifs are more likely to occur. Thus, one obtains a tool for analyzing future (or missing) graph structure, by being able to quantitatively compare different ways in which vertex sets may become (or already are) connected. Intuitively, we assume that a motif score should be high if the scores of participating edges are also high. This suggests one could reuse link prediction score functions.

Full extensive details of score functions, as well as more examples, are in the appendix.

A specific motif score function $s(M)$ will heavily depend on a targeted problem. In general, we define $s(M)$ as a function of V_M and E_M^* ; $s(M) = s(V_M, E_M^*)$. Here, $E_M^* = E_M \cup \overline{E}_{M,\mathcal{D}}$ are all the edges “that matter”: both edges in a motif (E_M) and the deal-breaker ones ($\overline{E}_{M,\mathcal{D}}$).

To obtain the exact form of $s(M)$, we harness existing link prediction scores for edges from E_M , when deriving $s(M)$ (details in § 3.4–§ 3.5). When using first-order link prediction methods (e.g., Jaccard), $s(M)$ depends on V_M and potential direct neighbors. With higher-order methods (e.g., Katz [49] or Adamic-Adar [2]), a larger part of the graph that is “around V_M ” is considered for computing $s(M)$. Here, our evaluation (details in Section 5) shows that, similarly to link prediction schemes [93], it is enough to consider a small part of G (1-2 hops away from V_M) to achieve high prediction accuracy for motifs.

However, we observe that simply extending link prediction fails to account for possible *correlations* between edges forming the motif (i.e., edges in E_M). Specifically, the appearance of some edges may impact (positively or negatively) the chances of one or more other edges in E_M . We provide score functions that consider such correlations in § 3.5.

3.4 Starting Simple: Motif Scores Based on Independent Links

There exist many score functions for link prediction [5, 6, 59, 76]. Similarly, one can develop motif prediction score functions with different applications in mind. As an example, we discuss score functions for a graph that models a set of people. An edge between two vertices indicates that two given persons know each other. For simplicity, let us first assume that there are no deal-breaker edges, thus $E_M^* = E_M$. For a set of people V_M , we set the score of a given specific motif $M = (V_M, E_M)$ to be the product of the scores of the associated edges: $s_{\perp}(M) = \prod_{e \in E_{M,\mathcal{N}}} s(e)$ where \perp denotes the independent aggregation scheme. Here, $s(e)$ is any link prediction score which outputs into $[0, 1]$ (e.g., Jaccard). Thus, also $s_{\perp}(M) \in [0, 1]$ by construction. Moreover, this score implicitly states that $\forall e \in E_{M,\mathcal{E}}$ we set $s(e) = 1$. Clearly, this does not impact the motif score $s_{\perp}(M)$ as the edges are already \mathcal{E} isting. Overall, we assume that a motif is more likely to appear if the edges that participate in that motif are also more likely. Now, when using the **Jaccard Score** for edges, the motif prediction score becomes $s_{\perp}(M)^J = \prod_{e_{u,v} \in E_{M,\mathcal{N}}} \frac{|N_u \cap N_v|}{|N_u \cup N_v|}$.

To **incorporate deal-breaker edges**, we generalize the motif score defined previously as $s_{\perp}^*(M) = \prod_{e \in E_M} s(e) \cdot \prod_{e \in \overline{E}_{M,\mathcal{D}}} (1 - s(e))$, where the product over E_M includes partial scores from the edges that belong to the motif, while the product over $\overline{E}_{M,\mathcal{D}}$ includes the scores from deal-breaker edges. Here, the larger the chance for a e to appear, the higher its score $s(e)$ is. Thus, whenever e is a deal-breaker, using $1 - s(e)$ has the desired diminishing effect on the final motif score $s_{\perp}^*(M)$.

3.5 Capturing Correlations between Links in a Motif

The main challenge is how to aggregate the link predictions taking into account the rich structural properties of motifs. Intuitively, using a plain product of scores implicitly assumes the independence of participating scores. However, arriving links may increase the chances of other links’ appearance in non-trivial ways. To capture such *positive correlations*, we propose heuristics based on the *convex linear combination of link scores*. To show that such schemes consider correlations, we first (Proposition 3.1) prove that the product P of any numbers in $[0, 1]$ is always bounded by the convex linear combination C of those numbers (the proof is in the appendix). Thus, our motif prediction scores based on the convex linear combination of link scores are always at least as large as the independent products of link scores (as we normalize them to be in $[0, 1]$, see § 3.6). The difference ($C - P$) is due to link correlations. Details are in § 3.5.1.

Proposition 3.1. *Let $\{x_1, \dots, x_n\}$ be any finite collection of elements from $U = \{x \in \mathbb{R} : 0 \leq x \leq 1\}$. Then, $\forall n \in \mathbb{N}$ we have $\prod_{i=1}^n x_i \leq \sum_{i=1}^n w_i x_i$, where $w_i \geq 0 \forall i \in \{1, \dots, n\}$ and subject to the constraint $\sum_{i=1}^n w_i = 1$.*

For *negative correlations* caused by deal-breaker edges, i.e., correlations that lower the overall chances of some motif to appear, we introduce appropriately normalized scores with a negative sign in the weighted score sum. The validity of this approach follows from Proposition 3.1 by noting that $\prod_{i=1}^n x_i \geq -\sum_{i=1}^n w_i x_i$ under the conditions specified in the proposition. This means

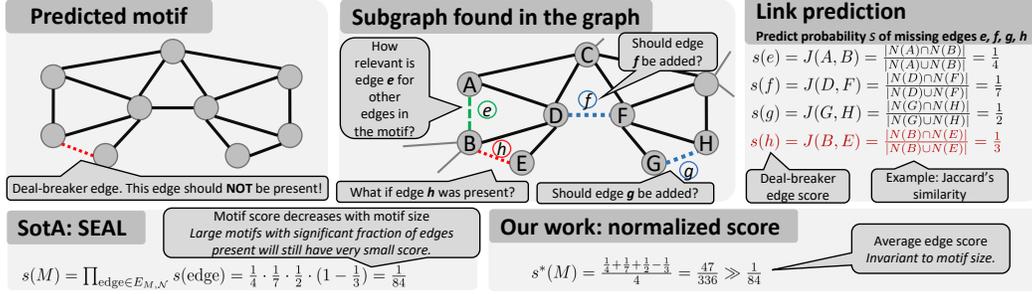


Figure 2: Illustration of the motif prediction problem, and example advantages of the heuristics proposed in this work, over existing methods such as SEAL link prediction.

that any combination of such negatives scores is always lower than the product of scores P ; the difference $|C - P|$ again indicates effects between links not captured by P . Details are in § 3.5.2.

3.5.1 Capturing Positive Correlation

In order to introduce positive correlation, we set the score of a given specific motif $M = (V_M, E_M)$ to be the convex linear combination of the vector of scores of the associated edges:

$$s(M) = f(s(\mathbf{e})) = \langle \mathbf{w}, s(\mathbf{e}) \rangle \quad (1)$$

Here, $f(s(\mathbf{e})) : [0, 1]^{|E_M|} \rightarrow [0, 1]$ with $|E_M| = |E_{V_M} \setminus \bar{E}_M|$ (i.e., not considering either \mathcal{I} ert or \mathcal{D} eat-breaker edges). In the weight vector $\mathbf{w} \in [0, 1]^{|E_M|}$, each component w_i is larger than zero, subject to the constraint $\sum_{i=1}^{|E_M|} w_i = 1$. Thus, $s(M)$ is a convex linear combination of the vector of link prediction scores $s(\mathbf{e})$. Finally, we assign a unit score for each existing edge $e \in E_{M, \mathcal{E}}$.

Now, to obtain a **correlated Jaccard score for motifs**, we set a score for each \mathcal{N} on-existing edge $e_{(u,v)}$ as $\frac{|N_u \cap N_v|}{|N_u \cup N_v|}$. \mathcal{E} xisting edges each receive scores 1. Finally, we set the weights as $\mathbf{w} = \frac{1}{|E_M|} \mathbf{1}$, assigning the same importance to each link within the motif M . This gives $s(M)^J = \frac{1}{|E_M|} \left(\sum_{e_{u,v} \in E_{M, \mathcal{N}}} \frac{|N_u \cap N_v|}{|N_u \cup N_v|} + |E_{M, \mathcal{E}}| \right)$. Any choice of $w_i > \frac{1}{|E_M|}$ places a larger weight on the i -th edge (and lower for others due to the constraint $\sum_{i=1}^{|E_M|} w_i = 1$). In this way we can incorporate domain knowledge for the motif of interest. For example, in the comparison of Figure 3, we set $\mathbf{w} = \frac{1}{|E_{M, \mathcal{N}}|} \mathbf{1}$ because of the relevant presence of \mathcal{E} xisting edges (each receiving a null score).

3.5.2 Capturing Negative Correlation from Deal-Breaker Links

To capture *negative correlation* potentially coming from deal-breaker edges, we assign negative signs to the respective link scores. Let $e \in E_M^* = E_M \cup \bar{E}_{M, \mathcal{D}}$. Then we set $s_i^*(e) = -s_i(e)$ if $e \in \bar{E}_{M, \mathcal{D}, \mathcal{N}}$, $\forall i \in \{1, \dots, |E_M^*|\}$. Moreover, if there is an edge $e \in \bar{E}_{M, \mathcal{D}, \mathcal{E}}$, we have $s^*(\mathbf{e}) = \mathbf{0}$. Assigning a negative link prediction score to a *potential Deal-breaker edge* lowers the score of the motif. Setting $s^*(\mathbf{e}) = \mathbf{0}$ when at least one *Deal-breaker edge* exists, allows us to rule out motifs which cannot arise. We now state a final motif prediction score:

$$s^*(M) = f(s^*(\mathbf{e})) = \max(0, \langle \mathbf{w}, s^*(\mathbf{e}) \rangle) \quad (2)$$

Here $s^*(M) : [0, 1]^{|E_M^*|} \rightarrow [0, 1]$ with $|E_M^*| \leq \binom{|V_M|}{2}$. Furthermore, we apply a rectifier on the convex linear combination of the transformed scores vector (i.e., $\langle \mathbf{w}, s^*(\mathbf{e}) \rangle$) with the rationale that any negative motif score implies the same impossibility of the motif to appear. All other score elements are identical to those in Eq. (1).

3.6 Normalization of Scores for Meaningful Comparisons and General Applicability

The motif scores defined so far consider only link prediction scores $s(e)$ with values in $[0, 1]$. Thus, popular heuristics such as Common Neighbors, Preferential Attachment, and the Adamic-Adar index do not fit into this framework. For this, we introduce a *normalized* score $s^{(e)}/c$ enforcing $c \geq \lceil \|\mathbf{s}(\mathbf{e})\|_\infty \rceil$ since the infinity norm of the vector of scores is the smallest value that ensures the desired mapping (the ceil function defines a proper generalization as $\lceil \|\mathbf{s}(\mathbf{e})\|_\infty \rceil = 1$ for, e.g., Jaccard [17]). To conclude, normalization also enables *the meaningful comparison of scores of different motifs which may differ in size or in their edge sets E_M* .

4 The SEAM Architecture for Training and Inference with GNNs

We argue that one could also *use neural networks to learn a heuristic for motif prediction*. Following recent work on link prediction [93, 95], we use a GNN for this; a GNN may be able to learn link correlations better than a simple hand-designed heuristic. Simultaneously, heuristics are still important as they do not require expensive training. We now describe the architecture of our design called SEAM (learning from Subgraphs, Embeddings and Attributes for Motif prediction).

4.1 Overview

Let $M = (V_M, E_M)$ be a motif to be predicted in G . First, we **extract the already existing instances of M in G** , denoted as $G_p = (V_p, E_p)$; $V_p \subseteq V$ and $E_p \subseteq E$. We use these instances G_p to **generate positive samples** for training and validation. To **generate negative samples** (details in § 4.2), we find subgraphs $G_n = (V_n, E_n)$ that do *not* form a motif M (i.e., $V_n = V_M$ and $E_M \not\subseteq E_n$ or $\overline{E_{M,\mathcal{D}}} \cap E_n \neq \emptyset$). Then, for each positive and negative sample, consisting of sets of vertices V_p and V_n , we **extract a “subgraph around this sample”**, $G_s = (V_s, E_s)$, with $V_p \subseteq V_s \subseteq V$ and $E_p \subseteq E_s \subseteq E$, or $V_n \subseteq V_s \subseteq V$ and $E_n \subseteq E_s \subseteq E$ (details in § 4.3). Here, we rely on the insights from SEAL [93] on their γ -decaying heuristic, i.e., it is G_s , the “surroundings” of a given sample (be it positive or negative), that are important in determining whether M appears or not. The nodes of these subgraphs are then **appropriately labeled** to encode the structural information (details in § 4.5). With these labeled subgraphs, we **train our GNN**, which classifies each subgraph depending on whether or not vertices V_p or V_n form the motif M . After training, we **evaluate the real world accuracy of our GNN** by using the validation dataset. The GNN architecture is detailed in § 4.6.

4.2 Positive and Negative Sampling

We need to provide a diverse set of samples to ensure that SEAM works reliably on a wide range of real data. For the positive samples, this is simple because the motif to be predicted (M) is specified. Negative samples are more challenging, because – for a given motif – there are many potential “false” motifs. In general, for each motif M , we generate negative samples using three strategies. (1) We first select positive samples and then remove a few vertices, replacing them with other nearby vertices (i.e., only a small number of motif edges are missing or only a small number of deal-breaker edges are added). Such negative samples closely resemble the positive ones. (2) We randomly sample V_M vertices from the graph; such negative samples are usually sparsely connected and do not resemble the positive ones. (3) We select a random vertex r into an empty set, and then we keep adding randomly selected vertices from the union over the neighborhoods of vertices already in the set, growing a subgraph until reaching the size of V_M ; such negative samples may resemble the positive ones to a certain degree. The final set of negative samples usually contains about 80% samples generated by strategy (1) and 10% each of samples generated by (2) and (3). This distribution could be adjusted based on domain knowledge of the input graph (we also experiment with other ratios). Strategies (2) and (3) are primarily used to avoid overfitting of our model.

As an example, let our motif M be a closed triangle. More precisely, $|V_M| = 3$ and $|E_M| = 3$. Consider a simple approach of generating negative samples, in which one randomly samples 3 vertex indices and verifies if there is a closed triangle between them. If we use these samples, in our evaluation for considered real world graphs, this leads to a distribution of 90% unconnected samples $|E_n| = 0$, 9% samples with $|E_n| = 1$ and only about 1% of samples with $|E_n| = 2$. Thus, if we train our GNN with this dataset, it would hardly learn the difference between open triangles $|E_n| = 2$ and closed triangles $|E_M| = 3$. Therefore, we provide our negative samples by ensuring that a third of

samples are open triangles $|E_n| = 2$ and another third of samples have one edge $|E_M| = 1$. For the remaining third of samples, we use the randomly generated vertex indices described above, which are mostly unconnected vertices $|E_M| = 0$.

For dense subgraphs, the sampling is less straightforward. Overall, the goal is to find samples with edge density being either close to, or far away from, the density threshold of a dense subgraph to be predicted. If the edge density of the sampled subgraph is lower than the density threshold it becomes a negative sample and vice versa. The samples chosen further away from the density threshold are used to prevent overfitting similar to strategies (2) and (3) from above. For this, we grow a set of vertices R (starting with a single random vertex), by iteratively adding selected neighbors of vertices in R such that we approach the desired density.

0 we start with a random vertex r , we insert it into a set R , and we iteratively test a given number c_1 of neighbors of vertices in R to find a vertex with many edges to the vertices in R , while having a degree of at least c_2 . By setting c_1 and c_2 appropriately, we can find samples that

Overall, we choose equally many positive and negative samples to ensure a balanced dataset. Furthermore, we limit the number of samples if there are too many, by taking a subset of the samples (selected uniformly at random). The positive and negative samples are split into a training dataset and a validation dataset. This split is typically done in a 9/1 ratio. To ensure an even distribution of all types of samples in these two datasets, we randomly permute the samples before splitting them.

4.3 Extracting Subgraphs Containing Positive and Negative Samples

To reduce the computational costs of our GNN, we do not use the entire graph G as input in training or validation. Instead, we rely on recent insights on link prediction with GNNs [93, 95], which illustrate that it suffices to provide a subgraph capturing the “close surroundings” (i.e., 1–2 hops away) of the vertices we want to predict a link between, cf. Section 2. We take an analogous assumption for motifs (our evaluation confirms the validity of the assumption). For this, we define the “surroundings” of a given motif $M = (V_M, E_M)$. Specifically, for a graph $G = (V, E)$ and the set of vertices $V_M \subseteq V$, the h -hop enclosing subgraph $G_{V_M}^h$ is given by the set of nodes $\{i \in V \mid \exists x \in V_M : d(i, x) \leq h\}$. To actually extract the subgraph, we simply traverse G starting from vertices in V_M , for h hops.

We then remove all dealbreaker edges from the subgraph of negative samples. For positive samples, we draw the existing edges $E_{M,\mathcal{E}}$ from the distribution of the negative samples; that is, we remove some edges from the subgraph such that it looks similar to a negative sample. Those removed edges are in essence the edges for which the GNN then has to correctly predict whether they will appear.

4.4 Node Embeddings

In certain cases, the h -hop enclosing subgraph might miss some relevant information about the motif in question. To alleviate this, while simultaneously avoiding sampling a subgraph with large h , we also generate a node embedding $X_E \in \mathbb{R}^{n \times f}$ which encodes the information about more distant graph regions using random walks. For this, we employ the established node2vec [40] with the parameters from DeepWalk [67]. f is the dimension of the low-dimensional vector representation of a node. We generate such a node embedding once and then only append the embedding vectors (corresponding to the nodes in the extracted subgraph) to the feature matrix of each extracted subgraph.

We obtain (cf. § 4.5)

$$X_s = (X_{s_i} \quad X_{s_E} \quad X_H \quad X_L \quad X_E) \in \mathbb{R}^{s \times (d+2f+2k)}.$$

Here, we also extend the SEAL approach called *negative injection* for more effective embeddings [93, 95]. The authors of SEAL observe that if embeddings are constructed using the edge set containing positive training samples, the GNN would focus on fitting this part of information. Thus, SEAL generates embedding based on the edge set containing *also* negative training samples, which ultimately improves accuracy. In SEAM, we analogously include all *potential* motif and deal-breaker edges E_M^* of all training samples to the input graph when generating the node embedding.

4.5 Node Labeling for Structural Features

In order to provide our GNN with as much structural information as possible, we introduce two node labeling schemes. These schemes serve as structural learning features, and we use them when constructing feature matrices of the extracted subgraphs, fed into a GNN. Let s be the total number of vertices in the extracted subgraph G_s and k be the number of vertices forming the motif. We call the vertices in the respective samples (V_p or V_n) the *inner* vertices since they form a motif sample. The rest of the nodes in the subgraph G_s are called *outer* vertices.

The first label is simply an enumeration of all the inner vertices. We call this label the *inner label*. It enables ordering each vertex according to its role in the motif. For example, to predict a k -star, we always assign the inner label 1 to the star central vertex. This inner node label gets translated into a one-hot matrix $H \in \mathbb{N}^{k \times k}$; $H_{ij} = 1$ means that the i -th vertex in V_M receives label j . In order to include H into the feature matrix of the subgraph, we concatenate H with a zero matrix $0_{(s-k)k} \in \mathbb{N}^{(s-k) \times k}$, obtaining $X_H = (H \ 0_{(s-k)k})^T$.

The second label is called the *outer* label. The label assigns to each outer vertex its distances to each inner vertex. Thus, each of the $s-k$ outer vertices get k labels. The first of these k labels describes the distance to the vertex with inner label 1. All these outer labels form a labeling matrix $L \in \mathbb{N}^{(s-k) \times k}$, appended with a zero matrix 0_{kk} , becoming $X_L = (0_{kk} \ L)^T \in \mathbb{N}^{s \times k}$. The final feature matrix X_s of the respective subgraph G_s consists of X_H , X_L , the subgraph node embedding matrix X_E and the subgraph input feature matrix $X_{s_i} \in \mathbb{R}^{s \times d}$, we have $X_s = (X_{s_i} \ X_E \ X_H \ X_L) \in \mathbb{R}^{s \times (d+f+2k)}$; d is the dimension of the input feature vectors and f is the dimension of the node embedding vectors.

4.6 Used Graph Neural Network Design

For our GNN model, we use the graph classification neural network DGCNN [94], used in SEAL [93, 95]. We now summarize its architecture. The first stage of this GNN consist of three graph convolution layers (GConv). Each layer distributes the vertex features of each vertex to its neighbors. Then, we feed the output of each of these GConv layers into a layer called k -sortpooling where all vertices are sorted based on their importance in the subgraph. After that, we apply a standard 1D convolution layer followed by a dense layer, followed by a softmax layer to get the prediction probabilities.

The input for our GNN model is the adjacency matrix of the selected h -hop enclosing subgraph $G_{V_s}^h$ together with the feature matrix X_s . With these inputs, we train our GNN model for 100 epochs. After each epoch, to validate the accuracy, we simply generate $G_{V_p}^h$ and $G_{V_n}^h$ as well as their feature matrix X_p and X_n from our samples in the validation dataset. We know for each set of vertices V_p or V_n , if they form the motif M . Thus, we can analyse the accuracy of our model by comparing the predictions with the original information about the motifs. Ultimately, we expect our model to predict the set of vertices V_p to form the motif M and the set of vertices V_n not to form the motif M .

5 Evaluation

We now illustrate the advantages of our correlated heuristics and of our learning architecture SEAM. We feature a representative set of results, extended results are in the appendix.

As **comparison targets**, we use motif prediction based on three link prediction heuristics (Jaccard, Common Neighbors, Adamic-Adar), and on the GNN based state-of-the-art SEAL link prediction scheme [93, 95]. Here, the motif score is derived using a product of link scores with no link correlation (“Mul”). We also consider our correlated heuristics, using link scores, where each score is assigned the same importance (“Avg”, $w = \frac{1}{|E_{M, \mathcal{N}}|}$), or the *smallest* link score is assigned the *highest* importance (“Min”). This gives a total of 12 targets. We then consider different variants of SEAM (e.g., with and without embeddings described in § 4.4).

As the **accuracy metric**, we use AUC (Area Under the Curve), the standard metric to evaluate the accuracy of any classification model in machine learning. We also consider a plain fraction of all correct predictions; these results resemble the AUC ones, see the appendix.

CN (Mul)	49.99 ± 0.45	49.78 ± 0.39	50.19 ± 0.47	50.13 ± 0.65	50.37 ± 0.79	51.55 ± 0.32	52.93 ± 0.61	55.42 ± 0.58	54.81 ± 0.58
CN (Min)	49.98 ± 0.33	49.72 ± 0.49	50.26 ± 0.59	50.35 ± 0.27	50.48 ± 0.32	51.77 ± 0.40	52.99 ± 0.75	54.75 ± 0.48	54.60 ± 0.85
CN (Avg)	49.76 ± 0.32	49.50 ± 0.64	50.18 ± 0.64	50.28 ± 0.51	50.91 ± 0.35	51.70 ± 0.59	53.32 ± 0.35	54.93 ± 0.77	54.20 ± 0.68
AA (Mul)	63.05 ± 0.71	62.09 ± 0.57	60.67 ± 0.95	54.95 ± 0.92	51.25 ± 0.63	51.40 ± 0.68	53.92 ± 0.52	55.15 ± 0.77	54.93 ± 0.61
AA (Min)	63.34 ± 0.68	62.81 ± 0.80	61.59 ± 0.94	54.81 ± 0.77	51.26 ± 0.38	51.60 ± 0.68	54.15 ± 0.89	54.59 ± 0.65	54.94 ± 0.27
AA (Avg)	63.96 ± 0.68	63.66 ± 0.48	62.71 ± 0.52	55.78 ± 0.74	51.28 ± 0.55	51.79 ± 0.62	54.52 ± 0.57	55.20 ± 0.53	54.55 ± 0.39
Jaccard (Mul)	67.17 ± 0.92	62.01 ± 0.72	59.71 ± 0.93	69.62 ± 1.09	57.60 ± 0.73	52.75 ± 0.97	51.75 ± 1.10	51.68 ± 0.85	50.93 ± 0.64
Jaccard (Min)	69.20 ± 0.80	67.11 ± 0.46	65.24 ± 0.80	73.88 ± 0.88	63.36 ± 1.17	56.50 ± 0.94	52.30 ± 0.54	51.86 ± 0.77	50.87 ± 0.53
Jaccard (Avg)	70.12 ± 0.78	68.59 ± 0.71	68.69 ± 0.77	75.35 ± 0.60	67.93 ± 0.87	61.22 ± 1.11	51.76 ± 0.91	49.74 ± 0.75	47.66 ± 0.58
SEAL (Mul)	76.68 ± 0.61	74.00 ± 0.50	71.80 ± 0.95	76.25 ± 1.90	63.66 ± 4.01	59.48 ± 4.87	68.53 ± 0.88	67.49 ± 1.27	67.88 ± 1.43
SEAL (Min)	77.15 ± 0.43	74.62 ± 0.55	73.11 ± 0.99	78.00 ± 1.49	69.70 ± 3.56	64.49 ± 5.47	66.40 ± 1.44	62.94 ± 1.98	62.88 ± 3.57
SEAL (Avg)	77.91 ± 0.91	75.98 ± 0.99	75.71 ± 0.66	77.50 ± 2.35	72.68 ± 3.21	66.95 ± 6.79	66.05 ± 0.78	65.14 ± 0.89	66.99 ± 1.32
SEAM, no embedding	86.24 ± 0.99	85.57 ± 0.94	88.61 ± 0.71	91.20 ± 1.03	96.16 ± 0.55	98.40 ± 0.22	83.39 ± 0.94	86.12 ± 0.66	87.86 ± 1.06
SEAM	90.78 ± 1.30	90.00 ± 1.84	91.53 ± 1.53	93.06 ± 0.61	97.26 ± 0.23	98.90 ± 0.18	83.81 ± 0.53	87.56 ± 0.79	88.59 ± 1.51
	3-star	5-star	7-star	3-clique	5-clique	7-clique	3-db-star	5-db-star	7-db-star

Figure 3: Comparison of different motif prediction schemes; SEAM is the proposed GNN based architecture. Other baselines use different link prediction schemes as building blocks; CN stands for Common Neighbors, AA stands for Adamic-Adar. We use the USAir graph, also used by the SEAL link prediction method [93, 95]. “ k -db-star” indicate motifs with deal-breaker edges considered.

5.1 Comparison of SEAM GNN, SEAL GNN, and Heuristics

We compare the accuracy of (1) our heuristics from Section 3, (2) a scheme using the SEAL link prediction, and (3) our proposed SEAM GNN architecture. The results are in Figure 3. First, we observe that the improvement in accuracy in SEAM almost always scales with the size of the motif. This shows that SEAM captures correlation between different edges (in larger motifs, there is more potential correlation between links). Importantly, the advantages and the capacity of SEAM to capture correlations, also hold in the presence of *deal-breaker edges* (“ k -db-star”). Here, we assign links connecting pairs of star outer vertices as deal-breakers (e.g., 7-db-star is a 7-star with 15 deal-breaker edges connecting its arms with one another). We observe that the accuracy for k -stars with deal-breaker edges is lower than that for standard k -stars. However, SEAM is still the best baseline since it appropriately learns such edges and their impact on the motif appearance.

Second, our correlated heuristics (“Avg”, “Min”) improve the motif prediction accuracy over methods that assume link independence (“Mul”). Overall, in heuristics, we observe that $AUC_{Mul} < AUC_{Min} < AUC_{Avg}$ (except for k -db-stars). This shows that different aggregation schemes have different capacity in capturing the rich correlation structure of motifs. In particular, notice that “Min” is by definition (cf. Proposition 3.1) a lower bound of the score $s(M)$ defined in § 3.5.1. This implies that it is the smallest form of correlation that we can include in our motif score given the convex linear combination function proposed in § 3.5.1.

Interestingly, the Common Neighbors heuristic performs poorly. This is due to the similar neighborhoods of the edges that have to be predicted. The high similarity of these neighborhoods is caused by our subgraph extraction strategy discussed in Section 4.3, where we select the existing motif edges of the positive samples in such a way as to mimic the edge structure of the negative samples. These results show also that different heuristics do not perform equally with respect to the task of motif prediction and further studies are needed in this direction.

The accuracy benefits of SEAM over the best competitor (SEAL using the “Avg” way to compose link prediction scores into motif prediction scores) range from 12% to almost 32%. This difference is even larger for other methods; it is because these comparison targets cannot effectively capture link correlations in motifs. This result shows that the edge correlation in motifs is important to accurately predict a motif’s appearance, and that it benefits greatly from being learned by a neural network.

Finally, we also provide the results for more arbitrary structures (clusters, communities), see Figure 4. Here, “ k -dense” indicates a cluster of k vertices, with at least 90% of all possible edges present. The results follow similar trends to those in Figure 3; SEAM outperforms all other baselines. Its accuracy also increase with the increasing motif size.

CN	83.44 ± 0.78	84.24 ± 0.61	84.74 ± 0.65
AA	83.05 ± 0.87	83.38 ± 0.78	82.15 ± 0.82
Jaccard	88.55 ± 0.37	86.44 ± 0.82	86.25 ± 0.42
SEAL	93.92 ± 1.91	92.66 ± 1.58	92.40 ± 0.82
SEAM, no embedding	98.16 ± 0.63	98.62 ± 0.36	99.45 ± 0.26
SEAM	98.86 ± 0.48	99.21 ± 0.28	99.66 ± 0.18
	11-dense	15-dense	19-dense

Figure 4: Comparison of prediction schemes as in Figure 3 for predicting dense subgraph motif described in § 4.2. All link prediction based schemes use the same motif score. We use the Yeast graph, also used by the SEAL link prediction method [93, 95].

5.2 Analysis of Different Aspects of Motif Prediction

We analyze the impact from embeddings. Interestingly, it consistently (by 0.2 – 4%) improves the accuracy while simultaneously *reducing the variance* in most cases by around 50% for cliques and dense clusters. We also consider other aspects, for example, we vary the number of \mathcal{E} existing edges, and even eliminate all such edges; the results follow similar patterns to those observed above. Finally, SEAM’s running times heavily depend on the used model parameters. A full SEAM execution on the Yeast graph dataset with 40,000 training samples and 100 training epochs does typically take between 15–75 minutes (depending heavily on the complexity of the motif, with stars and dense clusters being the fastest and slowest to process, respectively). The used hardware configuration includes an Intel 6130 @2.10GHz with 32 cores and an Nvidia V100 GPU; details are in the Appendix.

6 Related Work

We mention schemes related to prediction in Section 1. Next, many works exist on listing, counting, or finding different patterns [4, 6, 16, 20, 22, 24, 27, 28, 33, 34, 37, 43, 46, 54, 55, 57, 59, 68–70, 75, 81, 86]. Yet, they do not focus on prediction. Third, different works analyze the temporal aspects of motifs [48, 52, 53, 58, 66, 82], or use motif features for predicting links [1]. However, none of them considers prediction of general motifs. Moreover, there exists an extensive body of work on graph processing and algorithms, both static and dynamic (also called temporal, time-evolving, or streaming) [11–14, 18, 19, 25, 38, 50, 71, 74]. Still, they do not consider prediction of motifs. Finally, we use GNNs [73, 90, 97] for making predictions about motif appearance. An interesting venue of future work would be harnessing GNNs for other graph related tasks, such as compression [15, 21, 23].

There exist several recent frameworks for computations on GNNs [35, 42, 44, 45, 51, 56, 80, 85, 89, 92, 98]. While we use Pytorch Geometric, motif prediction can be integrated into any of these frameworks to enhance their processing capabilities. An interesting line of work would be to implement motif prediction using the serverless paradigm [7, 31, 47, 61], for example within one of recent dedicated serverless engines [79].

7 Conclusion & Discussion

Higher-order network analysis is an important approach for mining irregular data. Yet, it lacks methods and tools for predicting the evolution of the associated datasets. For this, we establish a problem of predicting general complex graph structures called motifs, such as cliques or stars. We illustrate its differences to simple link prediction, and then we propose heuristics for motif prediction that are invariant to the motif size and capture potential correlations between links forming a motif. Our analysis enables incorporating domain knowledge, and thus – similarly to link prediction – it can be a foundation for developing motif prediction schemes within specific domains.

While being fast, manual heuristics leave some space for improvements in prediction accuracy. To address this, we develop a graph neural network (GNN) architecture for predicting motifs. We show that it outperforms the state of the art, offering excellent accuracy, which *improves* with the growing size and complexity of the predicted motif. We also successfully apply our architecture to predicting more arbitrarily structured clusters, indicating its broader potential in mining irregular data.

Discussion: Societal Impact After extensively reviewing the available link prediction surveys [5, 6, 59, 76], we conclude that our work does not have any direct negative impact. Issues from more effective methods for mining data should be tackled using established privacy methods [3, 65, 77, 84].

Discussion: Limitations & Future Work We currently do not yet support certain motifs (e.g., k -clubs) and input graphs with weighted edges and vertices. Moreover, SEAM does not use distributed training for more scale. All these aspects are future work.

Acknowledgments and Disclosure of Funding

We thank Hussein Harake, Colin McMurtrie, Mark Klein, Angelo Mangili, and the whole CSCS team granting access to the Ault and Daint machines, and for their excellent technical support. We thank Timo Schneider for immense help with computing infrastructure at SPCL.

References

- [1] Ghadeer AbuOda, Gianmarco De Francisci Morales, and Ashraf Aboulnaga. 2019. Link prediction via higher-order motif features. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 412–429.
- [2] Lada A Adamic and Eytan Adar. 2003. Friends and neighbors on the web. *Social networks* 25, 3 (2003), 211–230.
- [3] Charu C Aggarwal and S Yu Philip. 2008. A general survey of privacy-preserving data mining models and algorithms. In *Privacy-preserving data mining*. Springer, 11–52.
- [4] Charu C Aggarwal and Haixun Wang. 2010. *Managing and mining graph data*. Vol. 40. Springer.
- [5] Mohammad Al Hasan, Vineet Chaoji, Saeed Salem, and Mohammed Zaki. 2006. Link prediction using supervised learning. In *SDM06: workshop on link analysis, counter-terrorism and security*.
- [6] Mohammad Al Hasan and Mohammed J Zaki. 2011. A survey of link prediction in social networks. In *Social network data analytics*. Springer, 243–275.
- [7] Ioana Baldini, Paul Castro, Kerry Chang, Perry Cheng, Stephen Fink, Vatche Ishakian, Nick Mitchell, Vinod Muthusamy, Rodric Rabbah, Aleksander Slominski, et al. 2017. Serverless computing: Current trends and open problems. In *Research Advances in Cloud Computing*. Springer, 1–20.
- [8] Vladimir Batagelj and Andrej Mrvar. 2006. Pajek datasets. (2006). <http://vlado.fmf.uni-lj.si/pub/networks/data/>.
- [9] Austin R Benson, Rediet Abebe, Michael T Schaub, Ali Jadbabaie, and Jon Kleinberg. 2018. Simplicial closure and higher-order link prediction. *Proceedings of the National Academy of Sciences* 115, 48 (2018), E11221–E11230.
- [10] Austin R Benson, David F Gleich, and Jure Leskovec. 2016. Higher-order organization of complex networks. *Science* 353, 6295 (2016), 163–166.
- [11] Maciej Besta, Armon Carigiet, Zur Vonarburg-Shmaria, Kacper Janda, Lukas Gianinazzi, and Torsten Hoefler. 2020. High-performance parallel graph coloring with strong guarantees on work, depth, and quality. *arXiv preprint arXiv:2008.11321* (2020).
- [12] Maciej Besta, Marc Fischer, Tal Ben-Nun, Dimitri Stanojevic, Johannes De Fine Licht, and Torsten Hoefler. 2020. Substream-Centric Maximum Matchings on FPGA. *ACM Transactions on Reconfigurable Technology and Systems (TRETS)* 13, 2 (2020), 1–33.
- [13] Maciej Besta, Marc Fischer, Vasiliki Kalavri, Michael Kapralov, and Torsten Hoefler. 2019. Practice of Streaming Processing of Dynamic Graphs: Concepts, Models, and Systems. *arXiv preprint arXiv:1912.12740* (2019).
- [14] Maciej Besta and Torsten Hoefler. 2015. Accelerating Irregular Computations with Hardware Transactional Memory and Active Messages. In *Proc. of the Intl. Symp. on High-Perf. Par. and Dist. Comp. (HPDC '15)*. 161–172.
- [15] Maciej Besta and Torsten Hoefler. 2018. Survey and taxonomy of lossless graph compression and space-efficient graph representations. *arXiv preprint arXiv:1806.01799* (2018).
- [16] Maciej Besta, Raghavendra Kanakagiri, Grzegorz Kwasniewski, Rachata Ausavarungnirun, Jakub Beránek, Konstantinos Kanellopoulos, Kacper Janda, Zur Vonarburg-Shmaria, Lukas Gianinazzi, Ioana Stefan, et al. 2021. SISA: Set-Centric Instruction Set Architecture for Graph Mining on Processing-in-Memory Systems. *arXiv preprint arXiv:2104.07582* (2021).
- [17] Maciej Besta, Raghavendra Kanakagiri, Harun Mustafa, Mikhail Karasikov, Gunnar Rätsch, Torsten Hoefler, and Edgar Solomonik. 2020. Communication-efficient jaccard similarity for high-performance distributed genome comparisons. In *2020 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 1122–1132.

- [18] Maciej Besta, Florian Marending, Edgar Solomonik, and Torsten Hoefler. 2017. SlimSell: A Vectorizable Graph Representation for Breadth-First Search. In *Parallel and Distributed Processing Symposium (IPDPS), 2017 IEEE International*. IEEE, 32–41.
- [19] Maciej Besta, Emanuel Peter, Robert Gerstenberger, Marc Fischer, Michał Podstawski, Claude Barthels, Gustavo Alonso, and Torsten Hoefler. 2019. Demystifying graph databases: Analysis and taxonomy of data organization, system designs, and graph queries. *arXiv preprint arXiv:1910.09017* (2019).
- [20] Maciej Besta, Michał Podstawski, Linus Groner, Edgar Solomonik, and Torsten Hoefler. 2017. To Push or To Pull: On Reducing Communication and Synchronization in Graph Computations. In *Proceedings of the 26th International Symposium on High-Performance Parallel and Distributed Computing*. ACM, 93–104.
- [21] Maciej Besta, Dimitri Stanojevic, Tijana Zivic, Jagpreet Singh, Maurice Hoerold, and Torsten Hoefler. 2018. Log (graph): a near-optimal high-performance graph representation. In *Proceedings of the 27th International Conference on Parallel Architectures and Compilation Techniques*. ACM, 7.
- [22] Maciej Besta, Zur Vonarburg-Shmaria, Yannick Schaffner, Leonardo Schwarz, Grzegorz Kwasniewski, Lukas Gianinazzi, Jakub Beranek, Kacper Janda, Tobias Holenstein, Sebastian Leisinger, et al. 2021. GraphMineSuite: Enabling High-Performance and Programmable Graph Mining Algorithms with Set Algebra. *arXiv preprint arXiv:2103.03653* (2021).
- [23] Maciej Besta, Simon Weber, Lukas Gianinazzi, Robert Gerstenberger, Andrey Ivanov, Yishai Oltchik, and Torsten Hoefler. 2019. Slim graph: Practical lossy graph compression for approximate graph processing, storage, and analytics. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–25.
- [24] Coen Bron and Joep Kerbosch. 1973. Algorithm 457: finding all cliques of an undirected graph. *Commun. ACM* 16, 9 (1973), 575–577.
- [25] Aydın Buluç and John R Gilbert. 2011. The Combinatorial BLAS: Design, implementation, and applications. *The International Journal of High Performance Computing Applications* 25, 4 (2011), 496–509.
- [26] Wenming Cao, Zhiyue Yan, Zhiquan He, and Zhihai He. 2020. A comprehensive survey on geometric deep learning. *IEEE Access* 8 (2020), 35929–35949.
- [27] Frédéric Cazals and Chinmay Karande. 2008. A note on the problem of reporting maximal cliques. *Theoretical Computer Science* 407, 1-3 (2008), 564–568.
- [28] Deepayan Chakrabarti and Christos Faloutsos. 2006. Graph mining: Laws, generators, and algorithms. *ACM computing surveys (CSUR)* 38, 1 (2006), 2.
- [29] Zhiqian Chen, Fanglan Chen, Lei Zhang, Taoran Ji, Kaiqun Fu, Liang Zhao, Feng Chen, and Chang-Tien Lu. 2020. Bridging the gap between spatial and spectral domains: A survey on graph neural networks. *arXiv preprint arXiv:2002.11867* (2020).
- [30] Diane J Cook and Lawrence B Holder. 2006. *Mining graph data*. John Wiley & Sons.
- [31] Marcin Copik, Grzegorz Kwasniewski, Maciej Besta, Michał Podstawski, and Torsten Hoefler. 2020. SeBS: A Serverless Benchmark Suite for Function-as-a-Service Computing. *arXiv preprint arXiv:2012.14132* (2020).
- [32] CSCS. 2021. Swiss national supercomputing center. (2021). <https://cscs.ch>.
- [33] Maximilien Danisch, Oana Balalau, and Mauro Sozio. 2018. Listing k-cliques in sparse real-world graphs. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 589–598.
- [34] David Eppstein, Maarten Löffler, and Darren Strash. 2010. Listing All Maximal Cliques in Sparse Graphs in Near-Optimal Time. In *Algorithms and Computation - 21st International Symposium, ISAAC 2010, Jeju Island, Korea, December 15-17, 2010, Proceedings, Part I*. 403–414. https://doi.org/10.1007/978-3-642-17517-6_36

- [35] Matthias Fey and Jan Eric Lenssen. 2019. Fast graph representation learning with PyTorch Geometric. *arXiv preprint arXiv:1903.02428* (2019).
- [36] Matthias Fey and Jan Eric Lenssen. 2019. Fast Graph Representation Learning with PyTorch Geometric. In *Proceedings of the International Conference on Learning Representations*, Vol. 7.
- [37] Brian Gallagher. 2006. Matching Structure and Semantics: A Survey on Graph-Based Pattern Matching.. In *AAAI Fall Symposium: Capturing and Using Patterns for Evidence Detection*. 45–53.
- [38] Lukas Gianinazzi, Pavel Kalvoda, Alessandro De Palma, Maciej Besta, and Torsten Hoefler. 2018. Communication-avoiding parallel minimum cuts and connected components. In *Proceedings of the 23rd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*. ACM, 219–232.
- [39] David Gibson, Ravi Kumar, and Andrew Tomkins. 2005. Discovering large dense subgraphs in massive graphs. In *Proceedings of the 31st international conference on Very large data bases*. 721–732.
- [40] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*. 855–864.
- [41] Tamás Horváth, Thomas Gärtner, and Stefan Wrobel. 2004. Cyclic pattern kernels for predictive graph mining. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 158–167.
- [42] Yuwei Hu, Zihao Ye, Minjie Wang, Jiali Yu, Da Zheng, Mu Li, Zheng Zhang, Zhiru Zhang, and Yida Wang. 2020. Featgraph: A flexible and efficient backend for graph neural network systems. *arXiv preprint arXiv:2008.11359* (2020).
- [43] Said Jabbour, Nizar Mhadhbi, Badran Raddaoui, and Lakhdar Sais. 2018. Pushing the Envelope in Overlapping Communities Detection. In *International Symposium on Intelligent Data Analysis*. Springer, 151–163.
- [44] Zhihao Jia, Sina Lin, Mingyu Gao, Matei Zaharia, and Alex Aiken. 2020. Improving the accuracy, scalability, and performance of graph neural networks with roc. *Proceedings of Machine Learning and Systems 2* (2020), 187–198.
- [45] Zhihao Jia, Sina Lin, Rex Ying, Jiaxuan You, Jure Leskovec, and Alex Aiken. 2020. Redundancy-Free Computation for Graph Neural Networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 997–1005.
- [46] Chuntao Jiang, Frans Coenen, and Michele Zito. 2013. A survey of frequent subgraph mining algorithms. *The Knowledge Engineering Review* 28, 1 (2013), 75–105.
- [47] Eric Jonas, Johann Schleier-Smith, Vikram Sreekanti, Chia-Che Tsai, Anurag Khandelwal, Qifan Pu, Vaishaal Shankar, Joao Carreira, Karl Krauth, Neeraja Yadwadkar, et al. 2019. Cloud programming simplified: A berkeley view on serverless computing. *arXiv preprint arXiv:1902.03383* (2019).
- [48] David Jurgens and Tsai-Ching Lu. 2012. Temporal motifs reveal the dynamics of editor interactions in Wikipedia. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 6.
- [49] Leo Katz. 1953. A new status index derived from sociometric analysis. *Psychometrika* 18, 1 (1953), 39–43.
- [50] Jeremy Kepner, Peter Aaltonen, David Bader, Aydin Buluç, Franz Franchetti, John Gilbert, Dylan Hutchison, Manoj Kumar, Andrew Lumsdaine, and Henning Meyerhenke. 2016. Mathematical foundations of the GraphBLAS. In *High Performance Extreme Computing Conference (HPEC), 2016 IEEE*. IEEE, 1–9.

- [51] Kevin Kinningham, Philip Levis, and Christopher Ré. 2020. GReTA: Hardware Optimized Graph Processing for GNNs. In *Proceedings of the Workshop on Resource-Constrained Machine Learning (ReCoML 2020)*.
- [52] Lauri Kovanen, Márton Karsai, Kimmo Kaski, János Kertész, and Jari Saramäki. 2011. Temporal motifs in time-dependent networks. *Journal of Statistical Mechanics: Theory and Experiment* 2011, 11 (2011), P11005.
- [53] Lauri Kovanen, Márton Karsai, Kimmo Kaski, János Kertész, and Jari Saramäki. 2013. Temporal motifs. In *Temporal networks*. Springer, 119–133.
- [54] Victor E Lee, Ning Ruan, Ruoming Jin, and Charu Aggarwal. 2010. A survey of algorithms for dense subgraph discovery. In *Managing and Mining Graph Data*. Springer, 303–336.
- [55] Elizabeth A Leicht, Petter Holme, and Mark EJ Newman. 2006. Vertex similarity in networks. *Physical Review E* 73, 2 (2006), 026120.
- [56] Shen Li, Yanli Zhao, Rohan Varma, Omkar Salpekar, Pieter Noordhuis, Teng Li, Adam Paszke, Jeff Smith, Brian Vaughan, Pritam Damania, et al. 2020. Pytorch distributed: Experiences on accelerating data parallel training. *arXiv preprint arXiv:2006.15704* (2020).
- [57] David Liben-Nowell and Jon Kleinberg. 2007. The link-prediction problem for social networks. *Journal of the American society for information science and technology* 58, 7 (2007), 1019–1031.
- [58] Paul Liu, Austin R Benson, and Moses Charikar. 2019. Sampling methods for counting temporal motifs. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*. 294–302.
- [59] Linyuan Lü and Tao Zhou. 2011. Link prediction in complex networks: A survey. *Physica A: statistical mechanics and its applications* 390, 6 (2011), 1150–1170.
- [60] Víctor Martínez, Fernando Berzal, and Juan-Carlos Cubero. 2016. A survey of link prediction in complex networks. *ACM computing surveys (CSUR)* 49, 4 (2016), 1–33.
- [61] Garrett McGrath and Paul R Brenner. 2017. Serverless computing: Design, implementation, and performance. In *2017 IEEE 37th International Conference on Distributed Computing Systems Workshops (ICDCSW)*. IEEE, 405–410.
- [62] Philipp Moritz, Robert Nishihara, Stephanie Wang, Alexey Tumanov, Richard Liaw, Eric Liang, Melih Elibol, Zongheng Yang, William Paul, Michael I Jordan, et al. 2017. Ray: A Distributed Framework for Emerging AI Applications. *arXiv preprint arXiv:1712.05889* (2017).
- [63] Huda Nassar, Austin R Benson, and David F Gleich. 2019. Pairwise link prediction. In *2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. IEEE, 386–393.
- [64] Huda Nassar, Austin R Benson, and David F Gleich. 2020. Neighborhood and PageRank methods for pairwise link prediction. *Social Network Analysis and Mining* 10, 1 (2020), 1–13.
- [65] Gayatri Nayak and Swagatika Devi. 2011. A survey on privacy preserving data mining: approaches and techniques. *International Journal of Engineering Science and Technology* 3, 3 (2011), 2127–2133.
- [66] Ashwin Paranjape, Austin R Benson, and Jure Leskovec. 2017. Motifs in temporal networks. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*. 601–610.
- [67] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 701–710.
- [68] T Ramraj and R Prabhakar. 2015. Frequent subgraph mining algorithms-a survey. *Procedia Computer Science* 47 (2015), 197–204.

- [69] Saif Ur Rehman, Asmat Ullah Khan, and Simon Fong. 2012. Graph mining: A survey of graph mining techniques. In *Seventh International Conference on Digital Information Management (ICDIM 2012)*. IEEE, 88–92.
- [70] Pedro Ribeiro, Pedro Paredes, Miguel EP Silva, David Aparicio, and Fernando Silva. 2019. A Survey on Subgraph Counting: Concepts, Algorithms and Applications to Network Motifs and Graphlets. *arXiv preprint arXiv:1910.13011* (2019).
- [71] Sherif Sakr, Angela Bonifati, Hannes Voigt, Alexandru Iosup, Khaled Ammar, Renzo Angles, Walid Aref, Marcelo Arenas, Maciej Besta, Peter A Boncz, et al. 2020. The Future is Big Graphs! A Community View on Graph Processing Systems. *arXiv preprint arXiv:2012.06171* (2020).
- [72] Ryoma Sato. 2020. A survey on the expressive power of graph neural networks. *arXiv preprint arXiv:2003.04078* (2020).
- [73] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. 2008. The graph neural network model. *IEEE transactions on neural networks* 20, 1 (2008), 61–80.
- [74] Edgar Solomonik, Maciej Besta, Flavio Vella, and Torsten Hoefler. 2017. Scaling betweenness centrality using communication-efficient sparse matrix multiplication. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. ACM, 47.
- [75] Lei Tang and Huan Liu. 2010. Graph mining applications to social network analysis. In *Managing and Mining Graph Data*. Springer, 487–513.
- [76] Ben Taskar, Ming-Fai Wong, Pieter Abbeel, and Daphne Koller. 2004. Link prediction in relational data. In *Advances in neural information processing systems*. 659–666.
- [77] Duygu Sinanc Terzi, Ramazan Terzi, and Seref Sagiroglu. 2015. A survey on security and privacy issues in big data. In *2015 10th International Conference for Internet Technology and Secured Transactions (ICITST)*. IEEE, 202–207.
- [78] Kiran K Thekumparampil, Chong Wang, Sewoong Oh, and Li-Jia Li. 2018. Attention-based graph neural network for semi-supervised learning. *arXiv preprint arXiv:1803.03735* (2018).
- [79] John Thorpe, Yifan Qiao, Jonathan Eyolfson, Shen Teng, Guanzhou Hu, Zhihao Jia, Jinliang Wei, Keval Vora, Ravi Netravali, Miryung Kim, et al. 2021. Dorylus: Affordable, Scalable, and Accurate GNN Training over Billion-Edge Graphs. *arXiv preprint arXiv:2105.11118* (2021).
- [80] Chao Tian, Lingxiao Ma, Zhi Yang, and Yafei Dai. 2020. PCGCN: Partition-Centric Processing for Accelerating Graph Convolutional Network. In *2020 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 936–945.
- [81] Etsuji Tomita, Akira Tanaka, and Haruhisa Takahashi. 2006. The worst-case time complexity for generating all maximal cliques and computational experiments. *Theor. Comput. Sci.* 363, 1 (2006), 28–42. <https://doi.org/10.1016/j.tcs.2006.06.015>
- [82] Sahar Torkamani and Volker Lohweg. 2017. Survey on time series motif discovery. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 7, 2 (2017), e1199.
- [83] Christian Von Mering, Roland Krause, Berend Snel, Michael Cornell, Stephen G Oliver, Stanley Fields, and Peer Bork. 2002. Comparative assessment of large-scale data sets of protein–protein interactions. *Nature* 417, 6887 (2002), 399–403.
- [84] Jian Wang, Yongcheng Luo, Yan Zhao, and Jiajin Le. 2009. A survey on privacy preserving data mining. In *2009 First International Workshop on Database Technology and Applications*. IEEE, 111–114.
- [85] Yuke Wang, Boyuan Feng, Gushu Li, Shuangchen Li, Lei Deng, Yuan Xie, and Yufei Ding. 2020. GNNAdvisor: An Efficient Runtime System for GNN Acceleration on GPUs. *arXiv preprint arXiv:2006.06608* (2020).

- [86] Takashi Washio and Hiroshi Motoda. 2003. State of the art of graph-based data mining. *Acm Sigkdd Explorations Newsletter* 5, 1 (2003), 59–68.
- [87] Duncan J Watts and Steven H Strogatz. 1998. Collective dynamics of ‘small-world’ networks. *nature* 393, 6684 (1998), 440–442.
- [88] Shiwen Wu, Fei Sun, Wentao Zhang, and Bin Cui. 2020. Graph neural networks in recommender systems: a survey. *arXiv preprint arXiv:2011.02260* (2020).
- [89] Yidi Wu, Kaihao Ma, Zhenkun Cai, Tatiana Jin, Boyang Li, Chenguang Zheng, James Cheng, and Fan Yu. 2021. Seastar: vertex-centric programming for graph neural networks. In *Proceedings of the Sixteenth European Conference on Computer Systems*. 359–375.
- [90] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. 2020. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems* (2020).
- [91] Chuxu Zhang, Dongjin Song, Chao Huang, Ananthram Swami, and Nitesh V Chawla. 2019. Heterogeneous graph neural network. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 793–803.
- [92] Dalong Zhang, Xin Huang, Ziqi Liu, Zhiyang Hu, Xianzheng Song, Zhibang Ge, Zhiqiang Zhang, Lin Wang, Jun Zhou, Yang Shuang, et al. 2020. Agl: a scalable system for industrial-purpose graph machine learning. *arXiv preprint arXiv:2003.02454* (2020).
- [93] Muhan Zhang and Yixin Chen. 2018. Link prediction based on graph neural networks. *arXiv preprint arXiv:1802.09691* (2018).
- [94] Muhan Zhang, Zhicheng Cui, Marion Neumann, and Yixin Chen. 2018. An end-to-end deep learning architecture for graph classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.
- [95] Muhan Zhang, Pan Li, Yinglong Xia, Kai Wang, and Long Jin. 2020. Revisiting Graph Neural Networks for Link Prediction. *arXiv preprint arXiv:2010.16103* (2020).
- [96] Ziwei Zhang, Peng Cui, and Wenwu Zhu. 2020. Deep learning on graphs: A survey. *IEEE Transactions on Knowledge and Data Engineering* (2020).
- [97] Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. 2020. Graph neural networks: A review of methods and applications. *AI Open* 1 (2020), 57–81.
- [98] Rong Zhu, Kun Zhao, Hongxia Yang, Wei Lin, Chang Zhou, Baole Ai, Yong Li, and Jingren Zhou. 2019. Aligraph: A comprehensive graph neural network platform. *arXiv preprint arXiv:1902.08730* (2019).

Appendix A: Proofs

We recall the statement of **Observation 1** in Section 3.1:

Consider vertices $v_1, \dots, v_k \in V$. Assuming no edges already connecting v_1, \dots, v_k , there are $2^{\binom{k}{2}} - 1$ motifs (with between 1 and $\binom{k}{2}$ edges) that can appear to connect v_1, \dots, v_k .

Proof. We denote as $E_k = \{\{i, j\} : i, j \in V_k \wedge i \neq j\}$ the edge set of the undirected subgraph (V_k, E_k) with $V_k \subseteq V$. The number of all possible edges between k vertices is $|E_k| = \binom{k}{2}$. Any subset of E_k , with the exception of the empty set, defines a motif. Thus the set of all possible subsets (i.e., the *power set* \mathcal{P}) of E_k is the set of motifs. Then, since $|\mathcal{P}(E_k)| = 2^{\binom{k}{2}}$, we subtract the empty set (which we consider as an invalid motif) from the total count to obtain the desired result. \square

We recall the statement of **Proposition 3.1** in Section 3.5:

Let $\{x_1, \dots, x_n\}$ be any finite collection of elements from $U = \{x \in \mathbb{R} : 0 \leq x \leq 1\}$. Then, $\forall n \in \mathbb{N}$ we have $\prod_{i=1}^n x_i \leq \sum_{i=1}^n w_i x_i$, where $w_i \geq 0 \forall i \in \{1, \dots, n\}$ and subject to the constraint $\sum_{i=1}^n w_i = 1$.

Proof. We start by noticing that $\prod_{i=1}^n x_i \leq \min\{x_1, \dots, x_n\}$. This is trivial to verify if $\exists x_i = 0$ for $i \in \{1, \dots, n\}$. Otherwise, it can be shown by contradiction: imagine that $\prod_{i=1}^n x_i > \min\{x_1, \dots, x_n\}$. We know that U is closed with respect to the product (i.e., $\prod_{i=1}^n x_i \in U \forall n \in \mathbb{N}$). Then, we can divide both sides by $\min\{x_1, \dots, x_n\}$, since we ruled out the division by zero, to obtain $\prod_{i=1}^{n-1} x_i > 1$. This implies $\prod_{i=1}^{n-1} x_i \notin U$, which contradicts that U is closed to the product. For the right side of the original statement, we know by definition that $x_i \geq \min\{x_1, \dots, x_n\} \forall i \in \{1, \dots, n\}$. Since $w_i \geq 0$, we can also write that $w_i x_i \geq w_i \min\{x_1, \dots, x_n\} \forall i \in \{1, \dots, n\}$. Thus, since U is an ordered set, we can state that $\sum_{i=1}^n w_i x_i \geq \sum_{i=1}^n w_i \min\{x_1, \dots, x_n\}$. But then, since $\sum_{i=1}^n w_i \min\{x_1, \dots, x_n\} = \min\{x_1, \dots, x_n\}$, we conclude that $\min\{x_1, \dots, x_n\} \leq \sum_{i=1}^n w_i x_i$. This ends the proof thanks to the transitive property. \square

Appendix B: Details of Datasets

In this section, we provide additional details on the various datasets that we used. We selected networks of different origins (biological, engineering, transportation), with different structural properties (different sparsities and skews in degree distributions).

USAir [8] is a graph with 332 vertices and 2,126 edges representing US cities and the airline connections between them. The vertex degrees range from 1 to 139 with an average degree of 12.8. Yeast [83] is a graph of protein-protein interactions in yeast with 2,375 vertices and 11,693 edges. The vertex degrees range from 1 to 118 with an average of 9.8. Power [87] is the electrical grid of the Western US with 4,941 vertices and 6,594 edges. The vertex degrees range from 1 to 19 with an average degree of 2.7.

Appendix C: Analysis and Decision Process of SEAM Model Parameters

We now discuss in more detail the selection of the SEAM model parameters.

Choosing learning rate and number of epochs

We first describe the process of tuning the hyperparameters for our motif prediction framework. In order to find the optimal learning rate for SEAM we try different learning rates as shown in Figures 5, 6 and 7. The associated hyperparameters are highly dependent on the specific motif to be predicted and on the used dataset. As an example, we analyze the hyperparameters for k -stars and k -cliques on the USAir graph dataset. The plots show that there is a sweet spot for the learning rate at 0.001-0.002. Any value below that rate is too small and our model cannot train its neural network effectively, while for the values above that, the model is unable to learn the often subtle differences between hard negative samples and positive samples. The number of epochs of the learning process can be chosen according to the available computational power of the user.

0.000125	80.89 ± 0.91	82.08 ± 0.64	85.35 ± 0.11	84.07 ± 0.95	85.08 ± 1.81	87.56 ± 1.18
0.00025	81.61 ± 0.82	84.68 ± 1.15	87.53 ± 0.36	84.70 ± 0.77	88.01 ± 1.76	91.17 ± 2.35
0.0005	83.80 ± 1.05	86.06 ± 0.42	89.55 ± 0.93	86.03 ± 0.97	91.31 ± 1.11	94.79 ± 1.02
0.001	84.50 ± 0.45	87.18 ± 0.62	90.83 ± 0.05	87.46 ± 0.98	93.64 ± 0.80	96.98 ± 0.53
0.002	86.26 ± 0.66	86.91 ± 0.93	90.42 ± 0.46	88.22 ± 0.53	94.77 ± 0.82	97.80 ± 0.78
0.004	86.92 ± 1.50	88.01 ± 2.25	88.04 ± 1.16	88.12 ± 1.30	94.76 ± 1.22	93.76 ± 8.11
0.008	83.87 ± 2.94	81.85 ± 4.97	84.32 ± 2.19	87.43 ± 2.45	81.89 ± 10.59	76.37 ± 11.53
0.016	71.03 ± 7.58	65.56 ± 5.81	63.58 ± 2.35	76.28 ± 8.69	66.94 ± 14.02	57.13 ± 5.09
0.032	54.14 ± 5.73	54.94 ± 9.88	50.00 ± 0.00	58.18 ± 9.06	56.33 ± 2.30	58.95 ± 7.42
	3-star	5-star	7-star	3-clique	5-clique	7-clique

Figure 5: AUC-Score comparison for different learning rates and a training duration of **50 epochs** on **USAir** graph.

SEAM parameters: proposed labels enabled, proposed embedding disabled, number of epochs = 50, training dataset size = 100,000

0.000125	82.46 ± 0.92	85.01 ± 0.56	88.56 ± 0.04	85.35 ± 0.82	87.99 ± 1.69	90.85 ± 0.64
0.00025	83.48 ± 0.59	86.83 ± 1.27	90.11 ± 0.20	86.28 ± 0.75	91.33 ± 1.28	94.61 ± 1.58
0.0005	85.62 ± 1.05	88.05 ± 0.21	91.53 ± 0.42	87.92 ± 0.83	94.75 ± 0.81	97.26 ± 0.47
0.001	86.58 ± 0.99	88.70 ± 0.57	92.27 ± 0.20	89.95 ± 1.34	95.77 ± 0.54	98.50 ± 0.37
0.002	88.40 ± 1.50	88.75 ± 0.96	91.79 ± 0.39	90.30 ± 0.43	96.78 ± 0.97	98.67 ± 0.31
0.004	89.68 ± 1.02	91.34 ± 2.97	90.93 ± 0.61	92.06 ± 2.70	96.62 ± 1.15	94.09 ± 8.94
0.008	86.72 ± 1.74	86.01 ± 2.49	88.13 ± 1.27	91.44 ± 3.06	90.45 ± 10.07	85.74 ± 8.56
0.016	75.29 ± 10.46	66.72 ± 7.92	63.58 ± 2.35	78.91 ± 11.09	70.96 ± 17.71	59.60 ± 11.00
0.032	55.32 ± 5.90	55.40 ± 10.80	50.00 ± 0.00	58.18 ± 9.06	56.25 ± 2.28	60.44 ± 6.98
	3-star	5-star	7-star	3-clique	5-clique	7-clique

Figure 6: AUC-Score comparison for different learning rates and a training duration of **100 epochs** on **USAir** graph.

SEAM parameters: proposed labels enabled, proposed embedding disabled, number of epochs = 100, training dataset size = 100,000

0.000125	83.50 ± 0.98	86.65 ± 0.29	90.12 ± 0.11	86.25 ± 0.72	90.06 ± 1.20	93.14 ± 0.58
0.00025	84.41 ± 0.75	87.66 ± 0.20	91.27 ± 0.41	87.46 ± 0.60	93.39 ± 1.13	96.42 ± 0.84
0.0005	86.47 ± 0.74	88.82 ± 0.39	92.23 ± 0.50	89.23 ± 0.60	95.96 ± 0.63	98.09 ± 0.27
0.001	88.00 ± 0.94	89.06 ± 0.47	92.72 ± 0.30	91.26 ± 1.25	96.46 ± 0.44	98.89 ± 0.33
0.002	89.43 ± 1.58	90.17 ± 1.52	92.37 ± 0.17	91.49 ± 1.43	97.26 ± 0.74	98.87 ± 0.31
0.004	91.52 ± 0.56	93.47 ± 2.44	91.30 ± 0.30	93.66 ± 3.01	97.33 ± 0.78	97.84 ± 1.23
0.008	87.46 ± 1.74	90.53 ± 3.21	90.39 ± 0.32	92.76 ± 2.78	93.93 ± 7.25	87.15 ± 9.33
0.016	76.81 ± 11.74	67.98 ± 10.33	63.58 ± 2.35	80.29 ± 11.41	72.57 ± 17.40	59.65 ± 10.99
0.032	55.32 ± 5.90	56.99 ± 12.11	50.00 ± 0.00	58.18 ± 9.05	56.25 ± 2.28	60.44 ± 6.98
	3-star	5-star	7-star	3-clique	5-clique	7-clique

Figure 7: AUC-Score comparison for different learning rates and a training duration of **150 epochs** on **USAir** graph.

SEAM parameters: proposed labels enabled, proposed embedding disabled, number of epochs = 150, training dataset size = 100,000

100	62.50 ± 12.72	54.80 ± 16.74	64.00 ± 8.23	70.90 ± 14.51	63.00 ± 11.07	62.90 ± 14.95	81.80 ± 8.26	81.20 ± 4.98	85.30 ± 6.90	74.30 ± 8.15	74.70 ± 6.81	72.00 ± 15.25
500	70.56 ± 3.98	70.10 ± 4.86	77.89 ± 3.59	75.64 ± 4.98	75.62 ± 4.15	74.07 ± 4.36	90.21 ± 1.54	88.08 ± 2.09	87.33 ± 3.65	74.54 ± 3.16	75.56 ± 4.06	75.45 ± 5.03
1k	75.77 ± 5.13	71.22 ± 3.73	79.67 ± 1.80	81.32 ± 3.42	80.09 ± 2.96	81.07 ± 4.66	91.75 ± 2.53	91.03 ± 2.60	89.70 ± 1.93	77.20 ± 2.79	77.10 ± 3.23	79.49 ± 3.40
5k	79.15 ± 2.13	80.03 ± 1.48	87.21 ± 0.66	87.37 ± 1.05	91.18 ± 1.73	96.05 ± 0.83	97.39 ± 0.72	97.17 ± 0.44	97.34 ± 0.26	81.54 ± 2.07	82.25 ± 2.00	84.28 ± 1.42
10k	82.45 ± 1.09	82.99 ± 0.80	89.35 ± 0.60	90.56 ± 0.85	93.87 ± 1.08	98.19 ± 0.50	98.04 ± 0.25	97.89 ± 0.28	98.06 ± 0.23	82.46 ± 1.43	85.36 ± 1.15	86.03 ± 1.06
25k	87.14 ± 1.55	87.63 ± 1.86	92.51 ± 0.79	90.22 ± 1.51	96.80 ± 0.69	98.74 ± 0.30	98.69 ± 0.15	98.81 ± 0.17	98.89 ± 0.18	83.29 ± 0.69	86.34 ± 0.53	88.69 ± 1.07
50k	87.41 ± 1.14	90.30 ± 2.35	93.80 ± 0.11	90.19 ± 0.93	96.46 ± 0.45	98.78 ± 0.46	99.16 ± 0.00	99.44 ± 0.00	99.32 ± 0.00	83.03 ± 0.79	86.49 ± 0.73	88.29 ± 0.90
	3-star	5-star	7-star	3-clique	5-clique	7-clique	11-dense	15-dense	19-dense	3-dbstar	5-dbstar	7-dbstar

Figure 8: AUC-Score comparison for different training dataset sizes on **USAir** graph

SEAM parameters: proposed labels enabled, proposed embedding enabled, learning rate = 0.002, number of epochs = 100

Analysis of different training dataset sizes

We also analyze the effect of different training dataset sizes on the prediction strength of SEAM. We want to assess the smallest number of samples that still ensures an effective learning process. Figure 8 shows the different accuracy results of SEAM, for different motifs and training dataset sizes. We observe that the accuracy strongly depends on the motif to be predicted. For example, a dense subgraph can be predicted with high accuracy with only 100 training samples. On the other hand, prediction accuracy of the 5-star motif improves proportionally to the amount of training samples while still requiring more samples (than plain dense subgraphs) for a high accuracy score. For all motifs, we set our minimal amount of training samples to 20,000 for positive and for negative ones.

Appendix D: Analysis of Different Variants of Motif Prediction in SEAM

Here, we analyze the effects and contributions from different variants of SEAM. First, we investigate the accuracy improvements due to our proposed labeling scheme in Section 4.5. Then, we empirically justify our approach to only sample the h -hop enclosing subgraph for small h (1–2). Finally, we evaluate the performance of every prediction method if there are no motif edges already present.

Effect on accuracy of our proposed labeling scheme

Figure 9 shows that our proposed labeling scheme generally has a positive impact on the accuracy of SEAM. The exception is the k -star motif. For $k = 3$, the labeling scheme significantly improves the accuracy. On the other hand, using $k > 3$ reduces the accuracy while simultaneously increasing the variance of test results. This effect can be explained with the implementation details of our labeling scheme. We remove every edges between all the motif vertices to calculate our k -dimensional distance labels. This procedure seems to misrepresent the structure of k -stars for $k > 3$. There are possible improvements to be gained in future work by further optimizing our labeling scheme.

Prediction Method	3-star	5-star	7-star	3-clique	5-clique	7-clique	3-dbstar	5-dbstar	7-dbstar
SEAM no labels	82.75 ± 0.75	94.71 ± 0.32	98.86 ± 0.12	89.21 ± 0.99	97.51 ± 0.35	98.88 ± 0.27	81.41 ± 0.95	85.39 ± 0.64	87.58 ± 0.64
SEAM	90.78 ± 1.30	90.00 ± 1.84	94.88 ± 2.28	93.06 ± 0.61	97.26 ± 0.23	98.90 ± 0.18	83.70 ± 0.82	87.56 ± 0.79	88.78 ± 1.49

Figure 9: Effect of our proposed labeling scheme on **USAir** graph.

SEAM parameters: h-hop = 1, proposed embedding enabled, learning rate = 0.002, number of epochs = 100, training dataset size = 100,000

Effect on accuracy of different h -hop enclosing subgraphs

Zhang et al. [93] motivated the use of small h -hop neighborhoods for their SEAL framework with the γ -decaying heuristic. We now provide additional data to backup this decision in SEAM. Figures 11 and 10 show that in most cases there is not much performance to be gained by sampling an h -hop enclosing subgraph with $h > 2$. This effect is especially striking for sparse graph datasets like the Power shown in Figure 11. The accuracy starts to drop significantly for $h > 2$. The only outlier in our little test was the 5-star motif shown in Figure 10. This effect was most likely caused by the specifics of this particular dataset and it does not reflect a trend for other graphs. An additional explanation could also be the non-optimal labeling implementation for the 5-star motif. These special cases do not justify to increase the neighborhood size of the motif in a general case.

Size of Neighborhood	3-star	5-star	7-star	3-clique	5-clique	7-clique
1-hop	86.20 ± 0.88	85.80 ± 0.93	88.61 ± 0.71	91.20 ± 1.03	96.16 ± 0.55	98.40 ± 0.22
2-hop	90.48 ± 0.76	92.80 ± 2.52	96.60 ± 0.66	92.41 ± 1.35	97.45 ± 0.19	98.97 ± 0.35
3-hop	91.46 ± 0.79	93.41 ± 1.59	95.87 ± 1.50	92.88 ± 1.21	97.59 ± 0.19	99.00 ± 0.32
4-hop	90.88 ± 1.52	95.58 ± 1.76	96.98 ± 0.70	93.21 ± 1.21	97.59 ± 0.52	99.04 ± 0.26

Figure 10: Comparison of different h -hop enclosing subgraphs used for our SEAM framework on **USAir** graph

SEAM parameters: proposed labels enabled, proposed embedding disabled, learning rate = 0.002, number of epochs = 100, training dataset size = 100,000

Size of Neighborhood	3-star	5-star	7-star	3-clique
1-hop	89.98 ± 1.28	92.72 ± 0.71	93.88 ± 0.71	72.19 ± 3.64
2-hop	90.18 ± 0.90	93.65 ± 0.62	94.90 ± 0.56	73.86 ± 3.61
3-hop	89.37 ± 1.03	90.47 ± 2.12	90.51 ± 4.94	73.42 ± 4.96

Figure 11: Comparison of different h -hop enclosing subgraphs used for our SEAM framework on **Power** graph

SEAM parameters: proposed labels enabled, proposed embedding disabled, learning rate = 0.002, number of epochs = 100, training dataset size = 100,000. **The Power graph does not contain enough 5-cliques and 7-cliques due to the sparsity of the graph.**

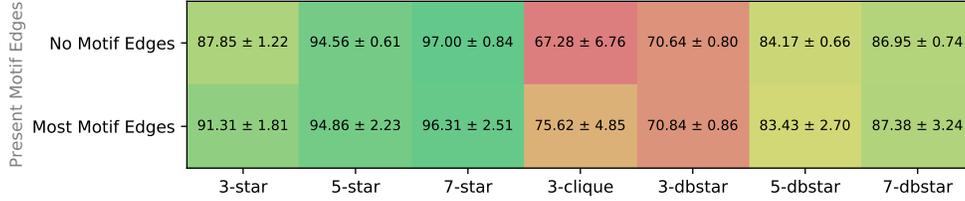


Figure 12: Comparison of the prediction accuracy of SEAM for different already present motif edges on **Power** graph

SEAM parameters: h-hop = 1, proposed labels enabled, proposed embedding enabled, learning rate = 0.002, number of epochs = 100 training dataset size = 100,000. **The Power graph does not contain enough 5-cliques and 7-cliques due to the sparsity of the graph.**

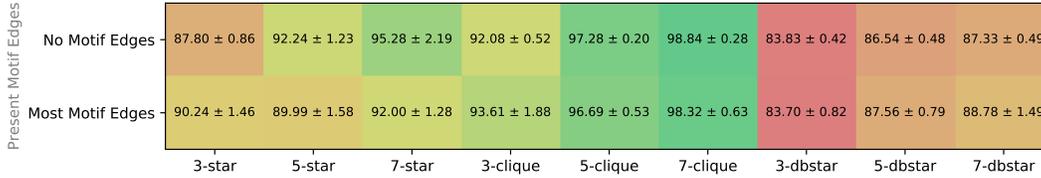


Figure 13: Comparison of the prediction accuracy of SEAM for different already present motif edges on **USAir** graph

SEAM parameters: h-hop = 1, proposed labels enabled, proposed embedding enabled, learning rate = 0.002, number of epochs = 100, training dataset size = 100,000

Effect on accuracy if no motif edges are previously known

We now illustrate that SEAM also ensures high accuracy when *no* or *very few* motif edges are already present, see Figures 12 and 13. Thus, we can conclude that SEAM’s prediction strength relies mostly on the structure of the neighborhood subgraph, embeddings, vertex attributes, and our proposed labeling scheme, and not necessarily on whether a given motif is already partially present. Outliers in this experiment are the 3-clique in the Power graph, the k -star motif with $k > 3$ in the USAir graph, and the 3-star motif in general. Still, there is no general tendency indicating that SEAM would profit greatly from the presence of most motif edges.

Appendix E: Analysis of Additional Datasets

We now analyze additional datasets, similarly to Section 5. An interesting result is the difference in accuracy for the Power graph dataset shown in Figure 14. This graph dataset is very sparse with very low average vertex degree. This result clearly shows very low accuracy of SEAL and other motif scores if there are just a few vertices in the neighborhood of the motif. The prediction accuracy for k -stars with deal-breaker edges is significantly better. This is caused by the properties of the positive samples discussed in Section 4.2. The prediction task of these positive samples boils down to predicting one motif edge, which has to be added, and several deal-breaker edges, that cannot appear. Due to the sparsity of the motif neighborhood, these deal-breaker edges are often predicted correctly to not appear, which significantly increases the prediction strength of SEAL and all the other motif scores.

In Figure 15, we use the Yeast graph dataset. The results in this dataset match the general trend of all our previous results. The exception in this dataset is the slight drop in accuracy for bigger stars and stars with deal-breaker edges. We conjecture that accuracy drop is caused by, for example (1) this dataset having a lot challenging negative samples or negative ones (4.2) for bigger motifs, (2) the neighborhoods of negative and positive samples being almost indistinguishable, or (3) limitations of our model.

Prediction Method	3-star	5-star	7-star	3-clique	3-dbstar	5-dbstar	7-dbstar
CN (Mul)	19.13 ± 0.29	25.72 ± 0.24	27.91 ± 0.24	51.11 ± 1.68	52.12 ± 0.40	50.28 ± 0.66	50.63 ± 0.94
CN (Min)	19.18 ± 0.22	25.62 ± 0.28	28.01 ± 0.22	51.11 ± 1.68	52.13 ± 0.49	50.28 ± 0.66	50.63 ± 0.94
CN (Avg)	19.27 ± 0.19	24.72 ± 0.35	26.15 ± 0.39	51.24 ± 1.72	53.14 ± 0.49	52.68 ± 0.69	52.93 ± 1.06
AA (Mul)	18.97 ± 0.50	26.26 ± 0.32	29.08 ± 0.41	42.04 ± 1.33	51.42 ± 0.49	50.35 ± 0.67	50.64 ± 0.94
AA (Min)	19.01 ± 0.22	25.95 ± 0.26	28.99 ± 0.49	42.06 ± 1.80	51.42 ± 0.49	50.35 ± 0.67	50.64 ± 0.94
AA (Avg)	19.20 ± 0.26	25.63 ± 0.34	27.59 ± 0.29	42.16 ± 2.44	52.33 ± 0.52	53.03 ± 0.73	53.50 ± 1.09
Jaccard (Mul)	20.47 ± 0.41	30.32 ± 0.26	33.59 ± 0.23	44.73 ± 2.70	50.76 ± 0.50	50.30 ± 0.67	50.62 ± 0.95
Jaccard (Min)	20.56 ± 0.21	30.62 ± 0.44	34.44 ± 0.58	47.27 ± 2.97	50.77 ± 0.50	50.30 ± 0.67	50.62 ± 0.95
Jaccard (Avg)	21.66 ± 0.49	31.30 ± 0.28	34.78 ± 0.34	48.17 ± 1.98	50.95 ± 0.55	51.04 ± 0.71	51.14 ± 0.94
SEAL (Mul)	25.90 ± 0.36	34.07 ± 0.38	37.05 ± 0.40	44.02 ± 2.21	45.61 ± 7.49	46.35 ± 6.54	47.46 ± 5.71
SEAL (Min)	24.58 ± 0.31	33.69 ± 0.20	36.92 ± 0.31	45.01 ± 2.79	47.53 ± 4.49	51.90 ± 2.29	54.40 ± 2.09
SEAL (Avg)	24.37 ± 0.23	33.51 ± 0.29	35.88 ± 0.59	45.48 ± 1.98	50.09 ± 3.00	50.26 ± 2.40	49.62 ± 3.43
SEAM, no embedding	89.98 ± 1.28	92.72 ± 0.71	93.88 ± 0.71	72.19 ± 3.64	70.34 ± 0.67	80.88 ± 1.06	84.28 ± 1.17
SEAM	92.64 ± 1.19	97.01 ± 0.46	98.74 ± 0.50	79.04 ± 3.21	71.34 ± 0.75	85.98 ± 0.72	90.47 ± 0.64

Figure 14: Comparison of different motif prediction schemes on **Power** graph.

SEAM is the proposed GNN based architecture. Other baselines use different link prediction schemes as building blocks; CN stands for Common Neighbors, AA stands for Adamic-Adar. “*k*-db-star” indicate motifs with deal-breaker edges considered. **The Power graph does not contain enough 5-cliques and 7-cliques due to the sparsity of the graph.**

Prediction Method	3-star	5-star	7-star	3-clique	5-clique	7-clique	3-dbstar	5-dbstar	7-dbstar
CN (Mul)	46.15 ± 0.54	44.26 ± 0.60	44.84 ± 0.68	50.80 ± 0.46	49.61 ± 0.46	50.10 ± 0.62	48.66 ± 0.70	50.69 ± 0.84	50.10 ± 0.53
CN (Min)	46.37 ± 0.79	43.96 ± 0.97	44.70 ± 0.46	50.77 ± 0.41	49.52 ± 0.72	50.02 ± 0.30	48.15 ± 0.53	50.73 ± 0.62	50.25 ± 0.73
CN (Avg)	46.27 ± 0.54	44.15 ± 0.99	44.36 ± 0.49	50.82 ± 0.45	49.24 ± 0.61	50.18 ± 0.77	48.10 ± 0.59	48.11 ± 0.53	47.18 ± 0.86
AA (Mul)	57.03 ± 0.73	54.50 ± 0.81	54.17 ± 0.92	54.44 ± 0.47	50.00 ± 0.77	50.50 ± 0.73	50.42 ± 1.03	50.44 ± 0.70	50.10 ± 0.69
AA (Min)	57.01 ± 0.81	55.15 ± 0.47	54.61 ± 0.75	54.26 ± 0.41	50.67 ± 0.63	50.45 ± 0.45	50.80 ± 0.58	50.42 ± 0.49	50.49 ± 0.57
AA (Avg)	57.76 ± 0.65	56.84 ± 1.03	56.67 ± 0.56	54.36 ± 0.62	50.26 ± 0.76	50.06 ± 0.75	51.23 ± 0.64	48.44 ± 1.09	48.25 ± 0.90
Jaccard (Mul)	57.49 ± 0.83	56.45 ± 0.51	56.34 ± 1.04	51.40 ± 0.73	50.58 ± 0.64	51.35 ± 0.63	49.67 ± 0.60	48.74 ± 0.64	48.81 ± 0.65
Jaccard (Min)	58.97 ± 0.64	60.18 ± 0.81	60.37 ± 0.96	53.18 ± 0.87	55.43 ± 1.10	54.35 ± 0.70	50.57 ± 0.56	48.88 ± 0.68	49.17 ± 0.57
Jaccard (Avg)	60.02 ± 0.86	62.18 ± 0.66	63.37 ± 0.98	54.37 ± 0.68	58.77 ± 0.69	60.43 ± 0.92	49.47 ± 0.64	46.59 ± 0.73	45.41 ± 0.69
SEAL (Mul)	71.82 ± 3.24	70.84 ± 1.09	69.59 ± 1.06	62.15 ± 4.01	59.66 ± 3.68	59.49 ± 1.69	62.85 ± 1.23	57.84 ± 1.28	55.12 ± 1.35
SEAL (Min)	72.94 ± 2.67	71.44 ± 1.30	69.68 ± 1.51	62.55 ± 3.77	61.89 ± 5.76	56.27 ± 2.72	60.23 ± 1.12	53.38 ± 1.24	53.46 ± 2.38
SEAL (Avg)	71.51 ± 1.63	72.13 ± 1.25	72.03 ± 0.91	66.26 ± 4.42	66.73 ± 4.74	61.72 ± 5.90	61.97 ± 1.42	57.98 ± 0.68	55.44 ± 0.84
SEAM, no embedding	89.81 ± 0.61	82.45 ± 0.87	82.28 ± 1.03	96.43 ± 0.36	95.74 ± 0.41	96.72 ± 0.23	84.42 ± 0.52	79.30 ± 0.98	79.83 ± 0.91
SEAM	90.13 ± 0.64	84.04 ± 1.21	83.69 ± 0.77	96.51 ± 0.25	96.90 ± 0.21	97.77 ± 0.31	84.37 ± 0.71	79.78 ± 0.66	81.54 ± 0.81

Figure 15: Comparison of different motif prediction schemes on **Yeast** graph.

SEAM is the proposed GNN based architecture. Other baselines use different link prediction schemes as building blocks; CN stands for Common Neighbors, AA stands for Adamic-Adar. “*k*-db-star” indicate motifs with deal-breaker edges considered.

Appendix F: Details of Implementation & Used Hardware

Our implementation² of SEAM and SEAL use the PyTorch Geometric Library [36]. We employ Ray [62] for distributed sampling and preprocessing, and RaySGD for distributed training and inference.

To run our experiments, we used the AULT cluster and the Piz Daint cluster at CSCS [32]. For smaller tasks, we used nodes from the AULT cluster such as AULT9/10 (64 AMD EPYC 7501 @ 2GHz processors, 512 GB memory and 4 Nvidia V100 GPUs), AULT23/24 (32 Intel Xeon 6130 @ 2.10GHz processors, 1.5TB memory and 4 Nvidia V100 GPUs), and AULT25 (128 AMD EPYC 7742 @ 2.25GHz processors, 512 GB memory and 4 Nvidia A100 GPUs). For larger, tasks we used our distributed implementation on the Piz Daint cluster (5704 compute nodes, each with 12 Intel Xeon E5-2690 v3 @ 2.60GHz processors, 64 GB memory and a Nvidia Tesla P100 GPU).

²Code will be available at http://spcl.inf.ethz.ch/Research/Parallel_Programming/motifs-GNNs/