

Sequential Motifs in Observed Walks

Timothy LaRock* Ingo Scholtes^{†‡} Tina Eliassi-Rad*[§]

July 11, 2022

Abstract

The structure of complex networks can be characterized by counting and analyzing network *motifs*. Motifs are small graph structures that occur repeatedly in a network, such as triangles or chains. Recent work has generalized motifs to temporal and dynamic network data. However, existing techniques do not generalize to *sequential* or *trajectory* data, which represents entities moving through the nodes of a network, such as passengers moving through transportation networks. The unit of observation in these data is fundamentally different, since we analyze observations of trajectories (e.g., a trip from airport A to airport C through airport B), rather than independent observations of edges or snapshots of graphs over time. In this work, we define *sequential motifs* in trajectory data, which are small, directed, and sequenced-ordered graphs corresponding to patterns in observed sequences. We draw a connection between counting and analysis of sequential motifs and Higher-Order Network (HON) models. We show that by mapping edges of a HON, specifically a k th-order DeBruijn graph, to sequential motifs, we can count and evaluate their importance in observed data. We test our methodology with two datasets: (1) passengers navigating an airport network and (2) people navigating the Wikipedia article network. We find that the most prevalent and important sequential motifs correspond to intuitive patterns of traversal in the real systems, and show empirically that the heterogeneity of edge weights in an observed higher-order DeBruijn graph has implications for the distributions of sequential motifs we expect to see across our null models.

1 Introduction

Motifs in complex networks are small graph structures that can be counted and analyzed from observed data. Methods for using motifs to characterize static

*Network Science Institute, Northeastern University, Boston, MA, USA

[†]Data Analytics Group, University of Zürich. Zürich, Switzerland

[‡]Chair of Machine Learning for Complex Networks, Center for Artificial Intelligence and Data Science (CAIDAS), Julius-Maximilians-Universität Würzburg, Germany

[§]Khoury College of Computer Sciences, Northeastern University, Boston, MA, USA

complex networks have been developed over the last 20 years [45, 4, 58, 73], and equivalent or similar concepts were investigated even earlier in the social sciences (see e.g., triadic analysis as reviewed in [74]). More recent research has focused on identifying motif structures in *temporal* or *dynamic* network data, which comes in the form of either timestamped independent edge observations or snapshots of a process over time [34, 33, 35, 77, 47, 71, 42]. In this work, we focus our attention on data that represents *sequences*, *trajectories*, or *walks* through networks.¹ Observations in this data are walks of varying lengths through a network. Each walk is observed independently and generally without temporal information about the time at which individual edges within the walk occurred. These walks may represent trajectories through physical networks (for example passengers or freight moving through a railway network) or through non-physical networks (for example sequences of proteins that make up the proteome of an organism, or sequences of words in a language such as n-grams [52]).

The goal of this work is to count and analyze *sequential motifs* in observed walks. Intuitively, a sequential motif is a small directed k -edge sequence-ordered multi-graph that is an abstract representation of a walk through a network. By sequence-ordered, we mean that each directed edge in the multi-graph is labeled with a unique integer $1 \leq j \leq k$, with these labels indicating the order in which the edges are traversed, e.g., the edge e_j is the j th edge to be traversed in the motif. In the simplest case, we can think of an edge from node A to node B, or the self-loop from node A to itself, as sequential motifs using 1 edge. These 1-edge structures are the objects of study in traditional network analysis, thus in practice we exclude them from our analyses; and in general, we will exclude self-loops. There are two sequential motifs that use 2 edges: the backtracking or bidirectional motif A-B-A (also called a 2-loop) and the chain motif A-B-C. These motifs are both relevant for understanding the interaction between a network structure and the processes that unfold on top of it. For example, in an airport network representing passenger movements, the sequential motif A-B-A represents a round-trip flight originating and terminating at the same airport.

In the left panel of Fig. 1, we show all simple (e.g., no multi-edges) directed motifs on 3-nodes. The left column shows motifs that are only interpretable in static network data, while the right column shows motifs that can be interpreted in either static or sequential contexts. A motif has a sequential interpretation if, starting from at least one node, there is a *directed Eulerian path*—a path that visits each edge exactly once—through the motif. Motifs without a Eulerian path starting from at least one node cannot possibly correspond to observed sequences of edges, since the lack of such a path indicates that there is at least 1 edge that cannot be traversed in sequence. In Figure 1, we have shown only motifs that are simple; however, the same property applies to motifs that are multi-graphs. For example, the motif A-B-A-B is itself a Eulerian path through the multi-graph with two labeled copies of the edge A-B and one edge B-A.

In the right panel of Figure 1, we show the observed frequency of the directed triangle sequential motif in data representing passenger flight itineraries

¹We will consider these terms interchangeable.

([70], see Section 4 for details). For air travel, the cycle motif corresponds to a trip with a direct flight in one direction and a layover in the other, or a trip that visits two different locations before returning to the origin, such as a multi-city business trip. The key takeaway from this comparison is that the directed triangle appears to be *overrepresented* compared to random expectation of sequential motifs. We say the triangle is overrepresented because it is observed considerably more often in the observed walk data than in walks randomly sampled uniformly from the edges of the complete DeBruijn graph, which corresponds to sampling from all possible 3-edge walks through the underlying directed network. However, based on static motifs computed from a directed and unweighted graph constructed from the flights (using the G-tries algorithm [56]), the motif appears to be underrepresented, as it occurs slightly less frequently than expected in randomizations of the underlying network structure. This sort of discrepancy is the motivation for our work.

Our methodology is tightly connected to recent research on Higher-Order Network Models (HONs). Here, higher-order refers to sequential and temporal correlations in how a network is traversed that violate the typical Markovian assumption implicit in the study of traditional, first-order networks [62, 61, 76, 60, 36, 38].² The violated assumption is that the $t + 1$ st step taken by a walker moving randomly on the edges of a network is dependent only on the position of the walker at time t , e.g., that the walker has no *memory* of where it was at times $i < t$. HONs for sequential data directly incorporate memory into a graphical representation, moving the Markovian assumption from edges between two nodes to walks through k nodes. These graphical representations can in turn be interpreted and analyzed using modified techniques from graph theory and network science.

In this paper, we count sequential motifs in observed data and evaluate their importance by mapping the edges of a particular HON, known as the *DeBruijn graph*, to their corresponding motif structures. In a DeBruijn graph of order k , each node represents a walk of length $k - 1$ (e.g., $A \rightarrow B$ for $k = 2$); and each directed, weighted edge represents the frequency of a length k walk through a traditional (or first-order) graph (e.g., $A \rightarrow B, B \rightarrow C$). Our work extends recent research showing the utility of DeBruijn graphs as representations for modeling correlations in observed walk data through networks to study motif structures [60, 38, 28].³

Definitions We compute sequential motifs from a dataset of n walks $\mathcal{S} = \{(s_1, w_1), (s_2, w_2), \dots, (s_n, w_n)\}$. Each $s_i = \langle u_1, u_2, \dots, u_{|s_i|} \rangle$ represents a walk on $(|s_i| - 1)$ edges through a directed graph $G = (V, E)$ with $|V|$ nodes and $|E|$ edges. For each walk s_i , the optional value w_i represents the *frequency* of

²This is in contrast to research on *higher-order structures* or *polyadic interactions*, such as hypergraphs or simplicial complexes [16, 6, 69].

³We note that our methodology can be applied to other HONs, for example the variable-order HON developed in [76]. We leave the investigation of motifs using these models for future work.

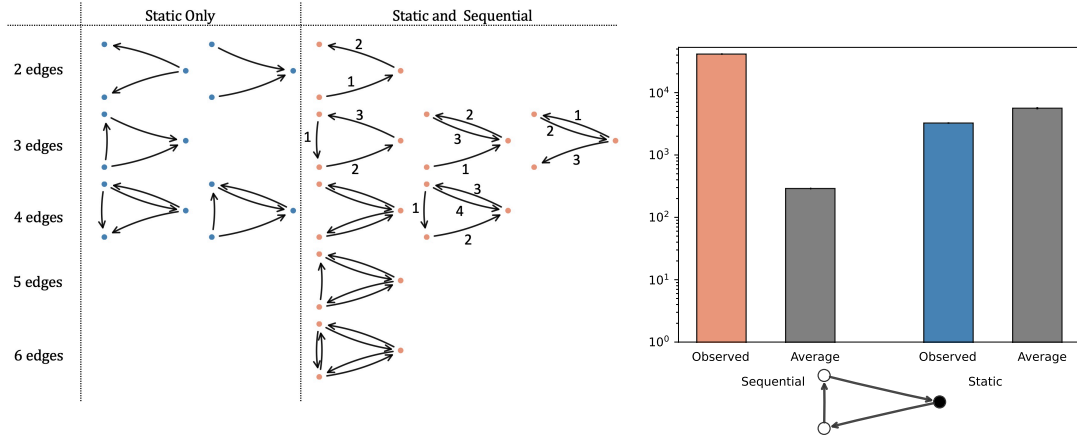


Figure 1: Left: Enumeration of 3-node simple directed motifs. Motifs in the right column have interpretations in both sequential data and static graphs, while those in the left column have no sequential interpretation because the edges involved cannot appear in sequence (equivalently there is no *Eulerian path* through the motif starting from any node). Right: Comparison of the frequency of the directed cycle motif on 3-nodes (i.e., a directed triangle) in data representing passenger trajectories through the airport-to-airport network. The first (orange) bar shows the observed frequency using the proposed sequential motif methodology. The second bar shows the same count but averaged over many randomizations of the data. The third (blue) bar shows the observed count in a static, directed, and unweighted graph, computed using the G-tries algorithm [56]. The last bar again shows the average frequency of the motif in randomized networks. Static motif analysis suggests cycles are under-represented, since the motif is more prevalent in randomized networks. In contrast, sequential motif analysis shows that the directed triangle is over-represented in the motifs compared with random realizations of the DeBruijn graph ensemble.

the observed walk in the data.⁴ If frequencies are not provided, \mathcal{S} is considered a multi-set where duplicate appearances of the same walk are aggregated and assigned a frequency equal to their multiplicity in \mathcal{S} . We denote as \mathcal{M}^k the set of k -edge sequential motifs, where each motif $m = \langle \sigma_1, \sigma_2, \dots, \sigma_{k+1} \rangle \in \mathcal{M}^k$ represents a sequence-ordered walk of length k edges through an alphabet, which is an ordered set of symbols $\Sigma = \{\sigma_1, \sigma_2, \dots, \sigma_\ell\}$, where $\ell \leq k+1$ is the maximum number of unique nodes in a k -edge walk. The alphabet may consist of any arbitrary set of symbols. The only requirement is that the number of unique symbols is at least as great as the number of unique nodes in any observed walk.

⁴For example, we may observe that 100 passengers bought the same round trip ticket $s_j = \langle u_1, u_2, u_1 \rangle$, and thus the observation s_j will have frequency $w_j = 100$.

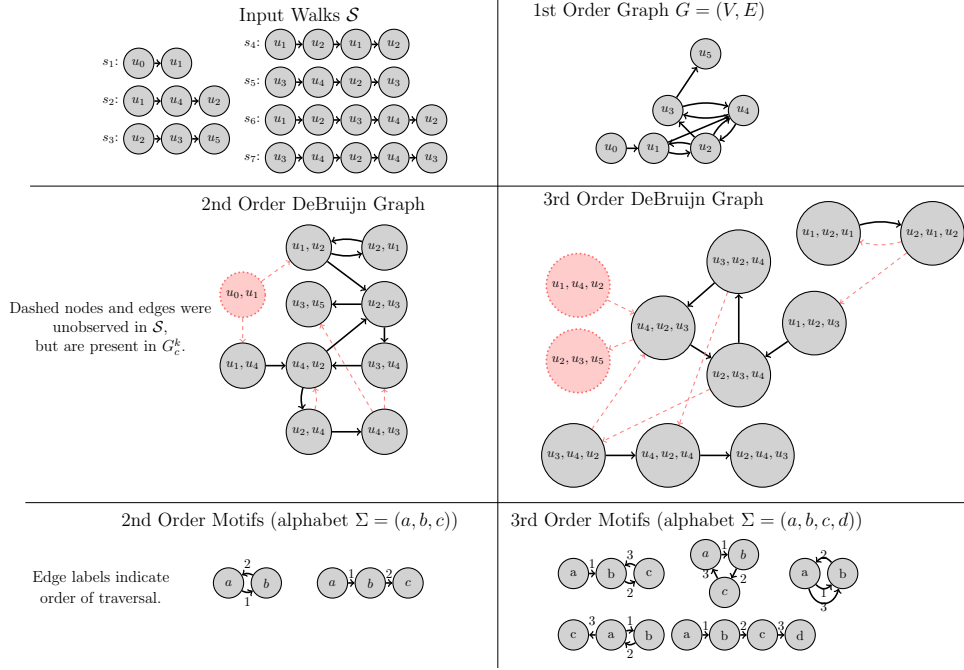


Figure 2: Example illustrating the relationship between walks, DeBruijn graphs, and sequential motifs. The top left panel shows the input dataset \mathcal{S} consisting of 7 walks on nodes $\{u_1, u_2, u_3, u_4, u_5\}$. The top right panel shows the first-order directed graph G constructed from the input data. The middle panels show DeBruijn graphs of order 2 and 3. Solid lines indicate the observed DeBruijn graph, while dashed lines indicate unobserved nodes and edges that exist in the complete DeBruijn graph G_c^k . The bottom panels show the sequential motifs counted from \mathcal{S} for orders 1 and 2. Labels on the edges of the motif indicate the order in which the edges are traversed.

Therefore, the worst-case alphabet size is the number of nodes N .⁵ Since we assume our input data are walks through a network, sequential motifs $m \in \mathcal{M}^k$ are sequence-ordered multi-graphs, meaning they may have parallel edges that represent multiple traversals between the same two nodes. Finally, we define a k th-order DeBruijn graph $G_k = (V_k, E_k)$ where V_k is the set of N_k k th-order nodes \vec{v} , each representing a walk of length $k - 1$ through G , and E_k is the set of M_k k th-order edges (\vec{v}, \vec{w}) , each representing a walk of length k . In a higher-order edge $(\vec{v}, \vec{w}) \in E_k$, the last $k - 1$ first-order nodes of \vec{v} must match the first $k - 1$ first-order nodes in \vec{w} , i.e., $\vec{v}[1, \dots, k] = \vec{w}[0, \dots, k - 1]$ for all

⁵In principle we can compute sequential motifs of any length k , provided that sequences of length at least k edges exist in the input data. However, larger values of k quickly become impractical to analyze. Therefore, we limit our results to motifs up to 3 edges, with the exception of Figure 1, where we show motifs up to 6-edges.

$$(\vec{v}, \vec{w}) \in E_k.$$

Contributions Our main contribution is to make explicit the connection between sequential motifs and DeBruijn graph models of observed walks. We use this connection to develop methods for evaluating the *importance* of a given sequential motif to a dataset or process (see Section 3.2). Since the edges of DeBruijn graphs correspond to walks of length k , we can interpret randomization of the DeBruijn graph structure as randomization of walks of length k (i.e., motif realizations, and vice versa). Thus our methods attempt to answer the question: *is this motif realized substantially more or less often in the observed data than expected at random?*

The key insight driving our approach is that we can count sequential motifs involving k edges by constructing a k th-order DeBruijn graph G_k and for every edge $e = (\vec{v}, \vec{w}) \in E_k$ that defines a k -edge walk $p_e = \vec{v} \cup \vec{w}[k-1]$ mapping the first-order nodes $u_i \in p_e, i \in 1 \dots k+1$ one-to-one with an alphabet Σ in sequential order. Projecting every edge in this way and counting the frequency of each projected motif pattern allows us to uncover patterns of traversal in observed walks. Fig. 2 shows a toy example illustrating this process. We then adopt and extend appropriate null models for DeBruijn graphs, which we use to evaluate the statistical importance of the motifs based on different randomizations of the walk data.

Our work makes the following contributions to the literature on understanding patterns in trajectories through complex networks:

- We define *sequential motifs*, which are small and directed sequence-ordered multi-graphs representing abstractions of walks through a directed graph.
- We show how counting sequential motifs can be achieved by mapping the edges of a DeBruijn graph to motifs and provide an algorithm called CAC for simultaneously constructing a DeBruijn graph and counting sequential motifs in time that scales linearly with the number of observed walks.
- We provide 3 methods for evaluating the importance of motifs based on sampling edges of DeBruijn graphs: sampling from the complete DeBruijn graph G_c^k , which corresponds to sampling uniformly from all possible walks through G ; sampling from the observed DeBruijn graph G^k , which corresponds to sampling uniformly from all observed walks (e.g., randomizing the frequency of the observed walks); and sampling edges based on an existing ensemble of weighted DeBruijn graphs known as HYPA [38], which corresponds to sampling walks of length k edges based on patterns of order $k-1$.
- We study sequential motifs in datasets of passenger trips through the domestic airport network in the United States, as well as clickstreams representing players navigating from random source pages to random targets in Wikipedia. We find a correspondence between how we expect people to

navigate the networks and the motifs that appear more or less often than expected at random based on our null models.

The remainder of the paper is organized as follows. In the next section, we review some work related to our research. In Section 3, we describe our methodology for computing sequential motifs and evaluating their importance. Then, in Section 4, we analyze empirical data using the proposed methods and discuss the circumstances in which sequential motifs are the right tool for analyzing a dataset. Finally, we conclude the paper in Section 5 with a discussion of future directions for this research.

2 Related Work

Motifs in network data have been studied in various forms across disciplines for decades [68, 55, 32]. The formulation presented in this work was inspired by work published mostly in the last twenty years, starting with the identification of network motifs as building blocks of complex networks [45]. This work measured and compared motifs across contexts, including gene regulatory networks (building on [65]), neuronal networks, food webs, electronic circuits, and the World Wide Web, using network randomization to define a notion of significance for each motif. Much of the early work on motifs in biological systems has been reviewed in [3] and more recently in [48]. The null models that underly motif comparisons are typically samples from *ensembles* of random graphs that preserve properties relevant to the particular networks under study [29]. Some work has also been done investigating *heterogeneous* network motifs, where heterogeneous refers to networks where nodes do not all have the same type [57].

The closely related concept of *closed frequent subgraph mining* was developed in parallel in the data mining community [78, 79]. In this formulation, the goal is to quickly identify interesting and maximal subgraph patterns in very large graph datasets, including over sets of graphs. These patterns are typically larger and more complex than what is studied in the literature on network motifs.

Another related literature addresses higher-order *structure* in graphs [6], for example by studying simplicial complexes [81, 49, 8, 30] and hypergraphs [7, 16], including work on hypergraph motifs [39]. Both of these areas are focused on simultaneous interaction among more than two nodes, sometimes referred to as polyadic interactions [16]. In contrast, our focus is on data that is characterized by dyadic interactions that happen in sequence.

Also related is the study of *subgraph frequencies*, most notably in [72]. This work defines a coordinate system for collections of disconnected graphs based on the frequency of small connected subgraphs, defined equivalently to motifs.

The work discussed so far largely addresses networks viewed as *static*, meaning that the nodes and edges remain the same over time. Yet another field of study has emerged to understand motifs in *temporal* network data [34, 33, 35, 77, 47, 71, 42]. Temporal network data comes in the form of timestamped edge observations $e = (u, v, t)$, where u and v are nodes and t is a timestamp.

Temporal, dynamic, or streaming network data is common, especially in (on-line) communication [1, 34, 33, 35, 77, 71] and also in biological systems [59]. The difference between temporal and sequential data is the unit of observation: temporal network data usually consists of independent and timestamped edge observations, while sequential data consists of observations of sequences of edges representing a walk through the network, usually without timestamps. In this data we begin from observations of sequences, but we do not necessarily know the order of the observations of the sequences themselves or the subobservations (e.g., individual edges) across different sequences, since we do not have timestamps. In contrast, in temporal network data we only have partial observations (individual edges) and need to infer or define a time-scale to find “whole” observations of walks, but we know for certain the order of the individual events given their timestamps. Although we do not address it in the current paper, this distinction does not necessarily impede us from analyzing temporal network data using sequential motifs. We can apply existing techniques to transform edge data into pathway, sequence, or observed walk datasets [50], then apply our sequential motif analysis to the resulting observed walk dataset. This moves the problem of determining the appropriate timescale at which to study the data to a pre-processing step [67].

Our work is also related to research on community detection using the line graph transformation, which is closely related to the DeBruijn graph [22]. An unweighted 2nd-order DeBruijn graph is the same as the line graph transformation of the first-order representation. In [22], the concept of *modularity* that is the basis of many community detection algorithms is generalized to line graphs to discover link partitions or communities in the network.

Our work is also related to the concept of k -motifs introduced by Sinatra et al. [66]. They identified sub-sequences of strings of symbols that they described as “fundamental units” of the process generating the strings, analogous to identifying words in English sentences where the punctuation has been removed. Although our methods seem similar at first glance, there is a fundamental difference in our respective deployments of the word “motif.” In our case, a motif is defined as an abstract sequence-ordered multi-graph, and trajectories through real networks can be viewed as realizations of a motif. Or, put another way, we are interested in motifs as *types* and walks as *tokens*, each of which is a realization of a type. In contrast, Sinatra et al [66] define a k -motif as a significantly re-occurring sub-sequence (e.g., a re-occurring token) in a sequential dataset. They defined and analyzed networks of k -motifs by first determining which sub-sequences were observed at higher than expected frequencies based on statistical significance of the observed frequencies in k th order Markov models. They then defined significant k -motifs as nodes in a motif graph, and connected pairs of motifs if they co-occurred in at least one sequence of the input data, where two motifs (i, j) co-occur if i appears immediately before j . Finally, they filtered out edges between k -motifs whose co-occurrences were not significant with respect to random expectation. They went on to show that community detection algorithms on k -motif networks can uncover patterns in diverse data, from clusters of proteins in the human proteome, to information cascades on

online social media platforms.

Recent work has attempted to understand the role of motifs in *dynamics* on networks [64, 59, 63]. Our work is related to that of Schwarze and Porter [63], who recently defined *process motifs* on graphs. A process motif is defined by the *walks* that are possible on a specific substructure of a network. Schwarze and Porter [63] show how the connection between structural and process motifs can be used to determine the role of different substructures in shaping a dynamical process on a network. Our work focuses on counting structural motifs from sequential data, but we note that because sequential data is the result of a discrete dynamical process (e.g., a random walk), our work blends structure and process, suggesting a close connection with this line of research.

We also build on the literature on DeBruijn graphs, which have been studied across fields for decades. In discrete mathematics, for example, the appearance and disappearance of cycles in DeBruijn graphs was a topic of interest in the 1970s [46, 41]. In biology and bioinformatics, variations of DeBruijn graphs are widely used to assemble DNA sequencing data into full genomes [51, 82, 31, 25]. In computer networks, optimal or near-optimal routing schemes have been found for DeBruijn graph representations [9, 43, 15, 24], and DeBruijn graphs have been used to design and analyze feedback shift registers in memory systems [40, 13]. Most recently, DeBruijn graphs have been introduced as an appropriate representation for sequential and pathway data on networks [60, 38, 28], which is the line of research we are contributing to most directly.

Finally, our research is related to the literature on network sampling. In general, network sampling is the process of gathering random samples of the network, for example random nodes or random edges (for a review, see [2]). These samples can then be used to approximate important network properties like the degree distribution [27, 80, 19, 53, 54, 17, 18]. Our work is related to the well-developed literature on sampling nodes uniformly at random using random walks [5, 14]. However, in our work we sample entire k -edge walks uniformly at random from all walks on k -edges, which is a fundamentally different problem. The work that is most closely related is about sampling induced subgraphs or graphlet realizations on a specific number of vertices [44, 11]. However, our work differs from these in that the motifs we are sampling are a subset of all possible motifs, since we only sample k -edge motifs that have valid sequential interpretations, and we focus exclusively on computing motifs from sequential data.

Our contribution is to link DeBruijn graphs as models of higher-order correlations in observed walk data to counting and analyzing sequential motifs. In Section 3, we describe our methodology for computing sequential motifs and evaluating their importance using DeBruijn graphs.

3 Proposed Methods

In this section we first present our method for simultaneously constructing a k th order DeBruijn graph and counting k -edge sequential motifs, which we call

CONSTRUCTANDCOUNT, or CAC. Then we describe a method for evaluating the importance of motifs using randomizations of walk datasets sampled from the edges of DeBruijn graphs.

3.1 Sequential Motif Counting

We start from a dataset of n trajectories $\mathcal{S} = \{(s_1, w_1), (s_2, w_2), \dots, (s_n, w_n)\}$ through the directed graph $G = (V, E)$ defined by the edges in \mathcal{S} . We are also given an alphabet $\Sigma = (\sigma_1, \sigma_2, \dots, \sigma_\ell)$ of arbitrary symbols (see problem statement in Section 1 for more details). Finally, we assume we have chosen an integer k that will be the length in edges of sequential motifs we are interested in computing.⁶

We provide the pseudocode for CAC in Algorithm 1. CAC takes as input a dataset of sequences \mathcal{S} and motif order k and proceeds as follows. First, CAC initializes the set of k th-order nodes V_k , edges E_k , and edge weights (equivalently frequencies) W_k to be empty. These will define the DeBruijn graph. CAC also initializes all entries of the lookup table $\mathcal{C}[m]$, indexed by motifs $m \in \mathcal{M}_k$, to 0. In practice, this initialization can be done dynamically the first time each motif m is found, or all $m \in \mathcal{M}_k$ can be given as input. This lookup table will contain the frequency counts of each motif. CAC then enters the outer for loop (line 3) that iterates over all of the input sequences and their associated weights: $(s, w) \in \mathcal{S}$. The default weight is 1. In the inner for loop (line 4), CAC slides a window of length k across s , where each index i corresponds to the first position of a k -edge subsequence seq starting at $s[i]$, computed in line 5. CAC then determines which motif the sequence seq corresponds to by projecting it using the `ProjectMotif(seq)` procedure, which maps each node $u \in seq$ to a distinct element σ_j in an alphabet Σ in order of appearance. For concreteness we can assume that `ProjectMotif(seq)` uses the ordered set of integers $\Sigma = 1, 2, \dots, k+1$ as its alphabet, but any ordered set of $k+1$ or more symbols will give equivalent results. After a node has been mapped to a symbol, it is replaced everywhere in the projected path with that symbol. Once every node in seq has been mapped, the procedure returns the sequential motif $m \in \mathcal{M}_k$ corresponding to seq . Now CAC increments the count of m in \mathcal{C} by the frequency of s , constructs the edge e from seq , and adds the node, edge and frequency to V_k , E_k , and W_k . Note that paths observed in \mathcal{S} with length shorter than k can not be included by definition, since $|s| - k$ will be negative and thus the inner for loop will not begin.

The runtime of CAC depends on three variables: (i) n , the number of sequences in \mathcal{S} ; (ii) $p(\ell)$, the distribution of lengths of sequences in \mathcal{S} ; and (iii)

⁶We do not offer a procedure for choosing a value of k automatically. We note that walks with fewer than k edges are ignored in our methodology, thus a value of k larger than the length of the longest trajectory in \mathcal{S} will yield no results. One could use statistical tools, such as those developed by Scholtes [60], to determine a k at which to analyze a dataset. One may also choose k based on the research question at hand; for example, a study of triangles must include $k = 3$. Finally, for exploratory analyses like those presented here, one may select values of k by starting with the smallest, $k = 2$, then studying motifs at increasing k as they see fit, or simply based on their intuition for the dataset or process.

Algorithm 1 CONSTRUCTANDCOUNT(S, k): Construct the k th order weighted DeBruijn Graph & count all length k motifs from sequence dataset S .

Input: $S = \{(p_1, w_1), (p_2, w_2) \dots, (p_N, w_N)\}$ (sequence dataset), k (order)

Output: $G_k \leftarrow (V_k, E_k, W_k)$ (DeBruijn graph), $\mathcal{C}[m]$ (motif counts)

```

1:  $V_k, E_k, W_k \leftarrow \emptyset$ 
2:  $\mathcal{C}[m] \leftarrow \emptyset \forall m \in \mathcal{M}_k$ 
3: for walk  $s \in S$  with weight  $w_s$  do
4:   for  $i$  from  $0, \dots, |s| - k$  do
5:      $\text{seq} = s[i, \dots, i + k]$ 
6:      $m \leftarrow \text{ProjectMotif}(\text{seq})$ 
7:      $\mathcal{C}[m] \leftarrow \mathcal{C}[m] + w_s$ 
8:      $\vec{u} \leftarrow \text{seq}[0, \dots, k - 1], \vec{v} \leftarrow \text{seq}[1, \dots, k]$ 
9:      $V_k \leftarrow V_k \cup \{\vec{u}, \vec{v}\}; E_k \leftarrow E_k \cup (\vec{u} \vec{v})$ 
10:     $W_k(\vec{u}, \vec{v}) = W_k(\vec{u}, \vec{v}) + w_s$ 
11: return  $G_k = (V_k, E_k, W_k), \mathcal{C}$ 

```

the motif length k . Assuming all paths are of the maximum length ℓ_{\max} , the worst case time is $O(n \cdot (\ell_{\max} - k))$. We note that in practice distributions of path lengths tend to have tails in large values, meaning that in real data the average path length ℓ_{avg} is likely to be much lower than the maximum (see Figure 4). Lastly, we note that the two loops in the CAC algorithm can be parallelized over the edges, since the operations for each edge (lines 4 through 10) are independent of any other edge (see Appendix A).

3.2 Measuring Motif Importance via Random DeBruijn Graphs

To evaluate the importance of a motif, we compare the count of that motif in the observed data with the average count of the same motif in many random samples from a null model. This is a flexible framework for determining importance, since the null model we choose determines which properties of the input data are randomized and to what extent, and therefore null models can be designed to test different hypotheses about the processes generating the observed data (similar to [45, 4]). Here we propose three null models. First, we describe a null model that relies on sampling from a *complete k th-order DeBruijn graph* G_k^c , where all possible k th-order edges based on the directed first-order topology G can be sampled with uniform probability. Second, we propose a null model that samples uniformly from all *observed k -edge walks* by sampling from the edges of the observed DeBruijn graph G^k . Finally, we adopt the null model proposed in [38], which samples from the edges of the observed k th-order DeBruijn graph based on a $k - 1$ st-order null model.

3.2.1 Sampling Uniformly at Random from All Possible k -edge Walks

Our goal for this null model is to find a uniform distribution for sequential motifs based on a given observed walk dataset \mathcal{S} . Thinking at the motif-distribution level, one way to do this would be to simply compare the observed distribution of motifs with the uniform distribution for a dataset of size $|\mathcal{S}|$, with the probability of each motif of a given length k being $\frac{|\mathcal{S}_k|}{|\mathcal{M}^k|}$, where here the quantity \mathcal{S}_k is the total number of k -edge observations in \mathcal{S} , and $|\mathcal{M}^k|$ is the number of k -edge motifs. However, since we are dealing with observed walks through networks, we know there is an underlying topology that constrains the motif distribution. As an example, consider the dataset \mathcal{S} containing only one observed walk, a chain on 2 edges $x - y - z$. In this case, the uniform distribution over all 2-edge motifs would be a strange choice, since the motif A-B-A would be given non-zero probability, despite being impossible based on the underlying topology (the edges $x - y$ and $y - z$). The lack of uniform reciprocity in real data means that these sorts of patterns in the first-order topology constrain the possible distribution of motifs.

Given this, we want a null model that accounts for the topological constraints imposed on the motif distribution by the underlying topology, while still being uniformly random with respect to the motif distribution. Our method achieves this by sampling edges uniformly from the complete DeBruijn graph G_k^c , which is equivalent to sampling uniformly from all possible walks through the first-order network topology.

Sampling from G_k^c is advantageous because it is conceptually simple to construct a complete k th-order DeBruijn graph G_k^c that includes every possible k th-order node and edge based on a directed first-order graph G . The procedure works as follows: first, take all of the edges in G and turn them into nodes in G_c^2 ; then, add edges between higher-order nodes (s, u) and (v, t) if $u = v$; and finally, repeat this procedure k times.⁷ Sampling an edge uniformly at random from the graph G_k^c is equivalent to sampling a k -edge walk uniformly at random from all possible k -edge walks through G .⁸ Importantly, this procedure is *not* equivalent to simulating random walks on the network topology. In Figure 3 we show a concrete case where sampling k -edge random walks would not result in a uniform distribution over all possible walks.⁹

In practice, computing complete DeBruijn graphs G_k^c for even relatively small values of k is computationally expensive, since the number of nodes and edges grows exponentially. In Appendix C, we give two strategies for sampling

⁷For further intuition, we note that the complete k th-order DeBruijn graph is equivalent to the line graph transformation applied k times to the first-order graph G .

⁸Note that our goal is *not* to sample nodes uniformly at random, which is a well-studied problem. In fact, we do not expect nodes to be sampled uniformly in this setting, since nodes with higher degree will be involved in more walks.

⁹We note that a Markovian walk on any k -regular graph, including a fully connected graph, is equivalent to sampling uniformly from all possible walks on that graph. This is because the set of possible walks starting from any node always has the same motif distribution, and the probability of sampling one of those possible walks is always equal, since the degree of every node is also equal.

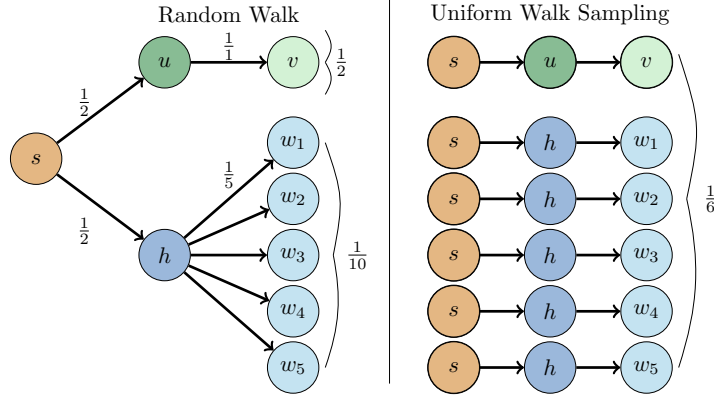


Figure 3: Sampling 2-edge random walks over the graph on the left is not the same as sampling uniformly from all possible 2-edge walks.

approximately uniformly from the edges of G_k^c . We also provide efficient code written in the Julia programming language [10] and building on the Graphs.jl package [23] for both exact and approximate procedures described in this paper on GitHub [37].

3.2.2 Sampling Uniformly at Random from Observed k -edge Walks

Sampling from all possible k -edge walks using G_k^c includes all possible walks through the first-order structure derived from the input data. However, there may be situations in which this null model is too permissive because some systems may have constraints that limit the possible k -edge walks. For example, in a railroad network there may be physical constraints, such as hard limits on where a train is able to reverse direction, or operational constraints, such as rules about how traffic congestion can be alleviated, that make some walks that are possible in the first-order topology impossible to observe at higher orders. In cases where such constraints exist, sampling walks from G_k^c may result in comparing motif distributions in data that does not make sense given the system constraints.

As a first step towards a more data-informed null model, we sample uniformly at random from the edges of the observed DeBruijn graph G^k . This is a uniform random sample over the observed k -edge walks, rather than all possible k -edge walks. If all of the edge weights in G^k are equal to 1, then our sampled motif distribution should be approximately equal to the observed motif distribution. However, since this is unlikely to be the case in many empirical datasets, we can study whether sampled counts for motifs corresponding to edges in G^k that have high weights are likely to be substantially smaller, and vice-versa, providing insight into whether individual motifs are important based on the topology, the process (e.g., the frequency of certain edges), or both.

3.2.3 Sampling k -edge Walks from HYP A

The previous two null models sample uniformly from a set of k -edge walks. We also evaluate motif importance by sampling from an ensemble derived from previous work on anomaly detection in DeBruijn graphs, which corresponds to sampling k -edge walks from a weighted $k - 1$ st-order model of the observed data [38]. HYP A represents a soft configuration model for the edge weights of a k th order observed DeBruijn graph, called the Generalized Hypergeometric Ensemble of Random Graphs [12]. We call the configuration model *soft* because samples from this ensemble preserve the average in and out weight for each node, not the exact degree sequence.

The HYP A sampling process works differently than the first two because we do not sample uniformly from all of the edges of the DeBruijn graph. In fact, we do not sample from the edges at all. Instead, HYP A defines a Hypergeometric distribution for the weight of each edge that takes into account the entire $(k - 1)$ st-order network structure [38]. Then, for each edge we sample an edge weight from its distribution, resulting in a k th-order DeBruijn graph with random edge weights that preserve the total weight and expected in- and out-strengths of the higher-order nodes. We then treat each edge weight as the multiplicity of the walk corresponding to that edge in our randomized walk dataset.¹⁰

We note that each of the null models we have proposed here can be customized to incorporate known correlations. For example, we may observe spatial correlations that indicate a high likelihood of certain walks occurring, and so rather than sampling uniformly, we may weight the walks based on a measure that captures this correlation. In fact, the Generalized Hypergeometric Ensemble that underlies HYP A is designed to incorporate this information using the propensity matrix [12]. We leave investigation of these bespoke null models for future work.

3.2.4 Discussion: Choosing Between Null Models

A natural question following the introduction of these three null models is: how should someone choose between them? Here we give some preliminary advice about how a researcher may go about deciding if one or the other is the right choice for their dataset. In general, we recommend simply comparing with all of the null models whenever possible, especially for more exploratory analyses. However, there are a few factors to consider in separating between the null models. We will organize these factors in descending order of statistical evidence.

In the case where there is evidence of statistical patterns at a certain order (walk length), then using the HYP A null model is likely a good choice. For example, if one applied methods like those in Scholtes 2017 [60] or LaRock 2020

¹⁰We believe that with some care it is possible to recast our procedure for sampling walks from HYP A as sampling uniformly from a multi-set of walks derived from the weighted $(k - 1)$ st-order DeBruijn graph constructed from \mathcal{S} . We leave this for future work or interested readers.

[38] and found statistical evidence for modeling the sequences at order $k=3$, then it might be a good idea to use the HYPA null model for 3-edge motifs.

If there is no evidence of a particular order at which to model the sequences, but there is good reason to believe that the observed data accurately reflects some external constraints of the system (as discussed in Section 3.2.2), then sampling from the observed DeBruijn graph may be a good choice. In this case, sampling from the observed DeBruijn graph G^k will preserve the system constraints while randomizing the frequency of the walks. However, determining whether such constraints exist or need to be modeled is ultimately up to the researcher.

If there is no evidence of a particular order at which to model the sequences, and the system is not inherently constrained to the observed walks, then sampling edges uniformly from the complete DeBruijn graph G_c^k may be the best choice of null model, since the only information encoded in this null model is the possible walks given the first-order structure. A second reason to use this null model is that it is an appropriate uniform null model for the distribution of sequential motifs that respects the network topology, as discussed in Section 3.2.1.

It is also worth noting that sampling from any of these null models entails constructing the observed DeBruijn graph G^k —either exactly, as in the HYPA, or as a subgraph, as in G_c^k . As a result, no matter which null model is chosen, one can always sample from the observed DeBruijn graph “for free”, i.e., without having to construct it separately.

Given this discussion, our advice is to start by determining whether the $(k-1)st$ -order patterns will be of interest. If so, then it would be wise to try HYPA. If there are concrete constraints on the observed walks based on properties of the system, sampling from G^k may be a good choice. Finally, in any case, but especially if there are no concrete constraints or higher-order patterns yet discovered, sampling from G_c^k is a reasonable choice.

4 Experiments

In this section, we present experiments on counting and analyzing sequential motifs in two datasets. We first describe the datasets—a transportation network and an online hyperlink navigation network—then we compare sequential motifs within and across the datasets. Appendix D contains more experimental results, and an implementation of our methodology is available online [37].

4.1 Data Description

Our first dataset consists of a large sample of flight itineraries representing trajectories of passengers through the domestic airport network in the United States in Quarter 1 of 2020 [70]. Each itinerary is a walk through the network corresponding to a starting airport, $i \geq 0$ layover airports, and a destination. We are also given frequencies for each walk corresponding to the number of people

who bought a ticket with that itinerary (e.g., a number w of people bought tickets with the itinerary JFK to Chicago O’Hare to Seattle, Washington).

The second dataset consists of walks through the Wikipedia network during successful runs of the Wikispeedia game, where a player was given a random target page in Wikipedia, then placed on a random start page and asked to navigate from the start to the target using only internal Wikipedia links [75]. These walks are not associated with frequencies, though we note that the DeBruijn graph edges are still weighted, since the same walks are repeated in different instances of the game. We present statistics of each of the datasets in Table 1, and show the distribution of sequence lengths in each dataset in Figure 4.

Table 1: Statistics of sequence datasets. Recall that $|V|$ is the number of nodes and $|E|$ is the number edges. ℓ_{avg} and ℓ_{max} are the average and maximum length of walks, respectively.

Dataset	$ \mathcal{S} $	ℓ_{avg}	ℓ_{max}	$ V $	$ E $
Flights	1,789,020	3.4	16	433	10,954
Wikispeedia	282,863	5.9	434	4,169	59,530

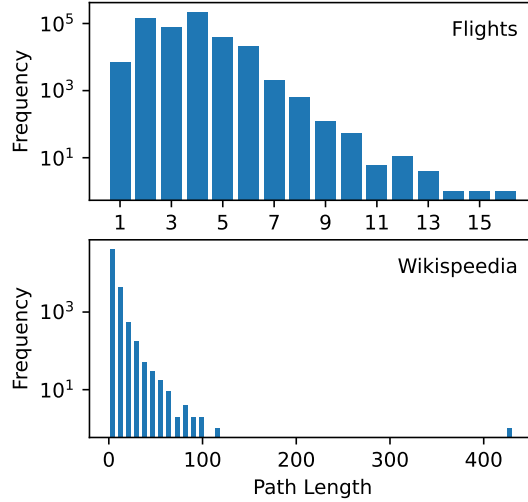


Figure 4: Histograms of sequence lengths in the passenger flights (top) and Wikispeedia (bottom) datasets. Both datasets have heavy tails in longer sequences, with most observations being relatively short and a small number of very long walks. Walks through the airport network are naturally shorter with a maximum of 10, while Wikispeedia walks have more variable length, with a relatively small number taking over 100 steps before finding the target page.

4.2 Sequential Motifs in Flight Patterns and Online Navigation

In Figure 5, we show the distribution of sequential motifs for orders 2 and 3 in both of our datasets. Each row of the figure corresponds to a single motif. The gray bar shows the observed frequency of each motif. The remaining bars show the average count of the motif over 10 randomized datasets from each of the following ensembles: uniformly from G_c^k ; uniformly from G^k ; the HYPA ensemble; unweighted first-order random walks; and weighted first-order random walks.¹¹ We will consider a motif *overrepresented* with respect to a null model if its observed frequency is greater than its frequency in samples from the null model, and *underrepresented* if its observed frequency is less than its frequency in the null model. If a motif has similar frequency in both the observed data and null model samples, we say its frequency is within expectation.

We begin with some general observations about the motif distributions. First, we note that all sampling methods tend to result in large numbers of chain motifs relative to all other motifs (second and last rows in Figure 5). This is not surprising since the number of possible chain motifs on k edges grows with the in- and out-degrees of the hubs, which are present in both networks. To illustrate this more concretely, consider a hub node h with degree d_h and a non-hub node u with degree $d_u \ll d_h$ that are connected in both directions. This structure creates only two 2-edge backtracking motifs ($u-h-u$ and $h-u-h$), but creates $d_h + d_u$ chain motifs ($u-h-*$ and $h-u-*$). Thus in any network with hubs (i.e., networks with heterogeneous degree distributions), we should expect our random samples to include more chains relative to other motifs.

A second observation is that simulating random walks is generally a good proxy for sampling uniformly from all possible walks via the edges of G_c^k (see Figure 3 for a reminder about why these are different processes). There is only one case where the motif count based on the random walker moving on the unweighted network is qualitatively different from G_c^k (Flights, 3rd row in Figure 5). There are a few possible explanations for this discrepancy. We believe the most likely explanation is that the random walker moving on the unweighted network is getting “caught” between two nodes that are only peripherally connected to the rest of the network. Since we do not see the same discrepancy for the random walker moving on the weighted network, it seems likely that when a walker takes edge weights into account to choose its next step there is a very small chance of it getting “caught” in this way, since the heaviest weights presumably lead towards the core of the network.

We now compare the distributions of each motif in turn, beginning from the first 2-edge motif at the top of Figure 5, which corresponds to the sequence A-B-A. In the flights data, paths that correspond to this motif represent round trips starting and ending at the same airport, whereas in the Wikispeedia game these correspond to a sort of “guess and check” behavior, where a player starts

¹¹We also computed standard deviations from the averages and found that they are typically two or more orders of magnitude smaller than the averages, making error bars too small to show on the figures.

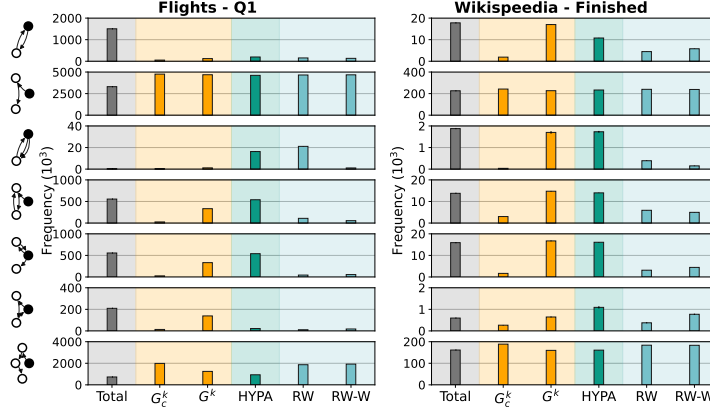


Figure 5: Comparing motif importance in US domestic flight paths (left) and paths through the Wikipedia page network in successful Wikispeedia games (right). Each row corresponds to a sequential motif (shown on the left). The gray panel shows the total number of motifs counted in the data. Note that the y-axes are measured in thousands. The orange bars show the average counts when sampling uniformly from the G_c^k and G^k ensembles; the dark green bar shows the average counts when sampling from HYPA and the blue bars show counts from unweighted (RW) and weighted (RW-W) random walks. A sequential motif is *overrepresented* with respect to a null model if its observed frequency is greater than its frequency in the null model, and *underrepresented* if its observed frequency is less than its frequency in the null model. In the Flights data, directed triangles are over-represented according to all ensembles. However, the results are mixed for the same motif in the Wikispeedia data, where it is considerably less prevalent. All null models except for the HYPA ensemble show the directed triangle as either within expectation or slightly underrepresented, while the HYPA ensemble suggests that the triangle is very underrepresented, appearing about twice as much in random samples. Since the directed triangle is a 3-edge sequential motif, the corresponding HYPA ensemble preserves the frequency of 2nd-order patterns in the Wikispeedia data. The fact that more triangles would be expected under this null model means that players opt not to close triangles as often as we would expect at random.

at page A, chooses a promising page B, then finds that the page does not contain any links that are better than those on the previous page, and so clicks a link back to page A. The motif is prevalent in both datasets, and is also overrepresented based on almost all ensembles. In the flights data, overrepresentation indicates that round trips occur substantially more often than expected in the real data than in any of the null models. This is not surprising, since people presumably prefer to take fewer flights, thus a direct flight in both directions

is ideal. However, it is worth noting that the results for 2-edge motifs in the flights data are essentially the same across all of the null models. To understand this, we first note that at $k = 2$, all of the null models, including HYPA, are sampling some variation of a 1st-order random walk. In fact, at $k = 2$ the HYPA ensemble and the weighted random walk (RW-W) are sampling from the same distribution—random walks on the weighted first-order topology—but with different strategies. In the flights data, regardless of which null model we choose, the result is the same: backtracks are very overrepresented, and chains are very underrepresented. One implication of this pattern is that the weights on the 2nd-order edges that correspond to backtracks must be non-uniform in the flights data, otherwise we would expect samples from G^k and HYPA to contain more backtracks. This makes intuitive sense because some round-trip flights—such as flights from major cities to Washington, D.C.—are likely to have extreme edge weights relative to less traveled trips. However, this does not hold in the Wikispeedia data, where sampling edges uniformly from the observed G^k results in an average count of backtrack motifs that is close to the observed count, and sampling from HYPA results in considerably more backtracks than the other null models. This indicates that the weight of the 2nd-order edges corresponding to the backtrack motif in the Wikispeedia data are relatively small on average, since sampling uniformly from those edges results in approximately the same number of backtracks. Since the Wikispeedia network is larger and more sparse than the flights network, it is not surprising that the distribution of 2nd-order backtrack weights is more uniform, as we show in the bottom left panel of Figure 6. Of course, we could have learned this information by examining the histograms in Figure 6 without the use of our null models; however, we discuss it here to draw a connection between the two frameworks.

The second 2-edge motif is the chain on 3 nodes: A-B-C. This motif is also prevalent in both datasets, and, as we observed above, all of our null models tend to sample many chains. In the flights data, the 2-edge chain appears underrepresented compared to all of the null models. The relatively low prevalence in the flights data is the flip side of the observation we made above: since people prefer to take direct flights in both directions, chains occur less often than expected at random. In the Wikispeedia data, the motif appears at approximately the same rates as all of the ensembles predict. In contrast to the flights, where the goal is often to return to the origin, in the Wikispeedia network the goal is to move away from the source and towards the target, so while backtracks are overrepresented, 2-edge chains are also an important part of the searching process.

The third motif from the top of Figure 5 is the first 3-edge motif, corresponding to a round trip that doubles back again (e.g., the sequence A-B-A-B). This motif is rarely observed in the flights dataset, and is observed at approximately the same rates as in samples from the G_c^k , G^k , and weighted random walk ensembles. However, the motif appears under-represented based on the HYPA and unweighted random walk ensembles. This is likely related to the prevalence of round trips in the observed data, which make the second backtracking edge much more likely than expected at random in a 2nd-order model or weighted

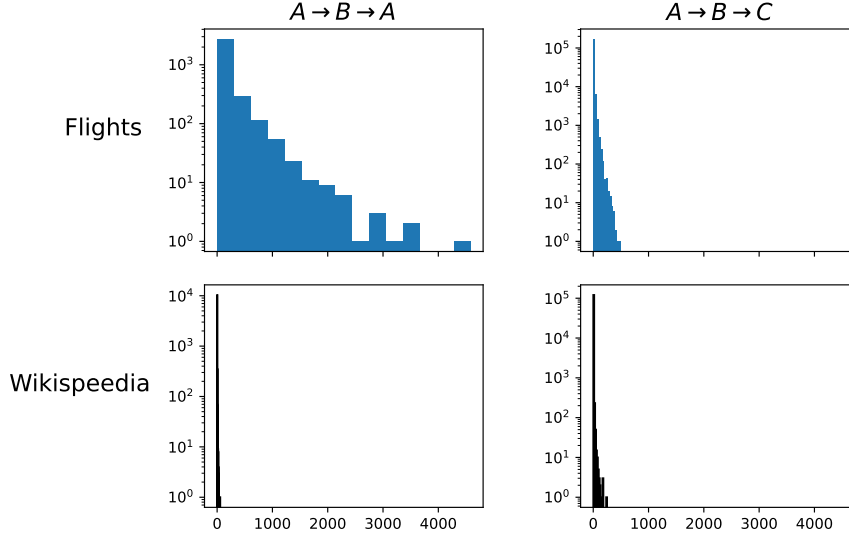


Figure 6: Histograms of weights for edges corresponding to the 2nd-order motifs in Flights (top) and Wikipedia (bottom) datasets. The distribution of edge weights for the backtrack motifs in the flights data (top left, A-B-A) is heavy tailed, while all others are not. This explains why all null models result in similar predictions for the flights backtrack motif.

random walk. In the Wikipedia data the motif appears to be within expectation for the observed G^k and HYPA ensembles, but is over-represented with respect to the others. This suggests that there are many observed k th-order edges corresponding to this motif, but the number of such edges is small compared to the total number of edges in G_c^k .

We will analyze the next two motifs (rows 4 and 5 from the top in Figure 5) together, since they are complementary to one another. The first represents the sequence A-B-A-C, while the second represents A-B-C-B. Taken together these motifs have an intuitive interpretation in the flights data: they represent a round trip with a layover at the same airport in each direction, broken into two motifs. In the first direction, we move from the source airport a_1 to a layover airport a_2 , then on to our destination to a_3 . On the return trip, we go from airport a_3 back to the layover airport a_2 , before finally returning again to a_1 . Therefore the whole trip is represented by the walk a_1, a_2, a_3, a_2, a_1 , which at $k = 3$ is broken into two motifs, the first starting from a_1 (A-B-C-B, 4th motif in Figure 5), and the second starting from a_2 (A-B-A-C, 5th motif). Therefore the 4th motif in Figure 5 corresponds to the first direction plus the edge (a_3, a_2) going back to the layover airport, while the 5th motif corresponds to the same edge plus the rest of the trip. These motifs are overrepresented with respect to the G_c^k and random walk ensembles in the flights data, which is intuitive since, as we noted in our discussion of the round trip motif, people tend to take

round trips when they fly, plus the fact that airlines send flights between the same airports on a regular schedule, making it likely that the layover will be the same in both directions. Again we see that the ensembles constrained to the observed walks result in counts similar to the observed frequencies. However, in this case the uniform samples from G^k still result in fewer observations of these motifs, while the counts based on HYPA are approximately the same as the observed data. The interpretation of these motifs in the Wikispeedia data is analogous if less intuitive, since they represent mediated round trips from one page, through two intermediary pages, then back where it started. This is indicative of a player giving up on a path and returning to an earlier page to try again, perhaps remembering that there was another promising link a few pages back. Accordingly, these motifs are also overrepresented in the Wikispeedia sequences with respect to the G_c^k and random walk ensembles, and just like in the flights data, the G^k and HYPA ensembles produce counts much closer to the observations.

Next we analyze the directed triangle (row 6), compared earlier to its static counterpart in Figure 1. In the flights data, this motif represents a trip with a direct flight in one direction and a layover in another, or a multi-city trip starting and ending at the same airport. In Wikispeedia, it corresponds to essentially the same discussion as in the last paragraph, but where the player did not need to backtrack because they found a link to the page to which they wanted to return. This motif is overrepresented with respect to all ensembles in the flights data, although the average count based on sampling uniformly from G^k is greater than the rest, suggesting that observed k th-order edges that correspond to triangles are a larger share of the edges in G^k compared to G_c^k . In the Wikispeedia data, where this motif occurs relatively rarely, all but the HYPA ensemble averages suggest the motif appears approximately within expectation, while the HYPA ensemble suggests the motif is underrepresented, indicating that the directed triangle is more likely to occur in k -edge walks through a 2nd-order model of the data.

Finally, we come to the 3-edge chain on 4 nodes (row 7). This is a very common motif in both datasets, and it appears to be underrepresented or within expectation according to all of the ensembles. Much of the same discussion of the 2-edge chain applies to the 3-edge chain.

Taken together with Figure 1, this analysis shows how sequential motifs can be used to understand a dataset in a way that is consistent with the inherent sequential nature of the process generating the data.

5 Conclusion

We introduced a method, CAC, for counting and analyzing *sequential motifs* on k edges in observed walk data using DeBruijn graph ensembles and random walks. A major advantage of sequential motifs is that their interpretation is straightforward: we count the prevalence of a particular structure *in the data*, preserving the sequential information that would be lost by aggregating to a

first-order network, the first step of a more traditional approach. We presented three sampling procedures to assist in evaluating the importance of sequential motifs: (1) a uniform k -edge walk sampling procedure from the edges of a complete DeBruijn graph G_c^k , (2) a uniform sampling procedure over observed walks based on the edges of the observed DeBruijn graph G^k , and (3) a sampling procedure based on the Generalized Hypergeometric Ensemble of DeBruijn Graphs defined in HYPA [38]. We showed that sequential motifs are substantially different than static network motifs, and analyzed two datasets, highlighting the subtleties between the various null models and comparing to simply sampling random walks on the 1st-order graph.

There are numerous future directions for this research. For example, we can design bespoke DeBruijn graph null models for specific datasets based on external correlations such as geographical or conceptual distance between nodes. The convergence of our sampling methods to the true motif distribution, as well as the related question of when a random walk ensemble is good enough, also remain open. It also seems likely that random walk simulations could be designed that carefully account for node degrees, resulting in sets of walks that are sampled uniformly from all possible walks.

Funding IS acknowledges financial support by the Swiss National Science Foundation through grant 176938. TL and TER were funded by the Combat Capabilities Development Command Army Research Laboratory through Cooperative Agreement W911NF-13-2-0045 and by the Under Secretary of Defense for Research and Engineering under Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Combat Capabilities Development Command Army Research Laboratory, Under Secretary of Defense for Research and Engineering, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes not withstanding any copyright notation hereon.

Acknowledgements TL thanks Vahan Nanumyan for preliminary discussions and analyzes that led to this work; Leo Torres for advice and discussion about the methodology; and Brennan Klein for help with designing the visualizations.

References

- [1] J. Abello, T. Eliassi-Rad, and N. Devanur. Detecting Novel Discrepancies in Communication Networks. In *Proceedings of the 10th IEEE International Conference on Data Mining, ICDM'10*, pages 8–17, 2010.
- [2] N. K. Ahmed, J. Neville, and R. Kompella. Network Sampling: From Static to Streaming Graphs. *ACM Transactions on Knowledge Discovery from Data*, 8(2):7:1–7:56, 2014.
- [3] U. Alon. Network motifs: Theory and experimental approaches. *Nature Reviews Genetics*, 8:450, June 2007.

- [4] Y. Artzy-Randrup, S. J. Fleishman, N. Ben-Tal, and L. Stone. Comment on “Network Motifs: Simple Building Blocks of Complex Networks” and “Superfamilies of Evolved and Designed Networks”. *Science*, 305(5687):1107–1107, Aug. 2004.
- [5] B. A. Bash, J. W. Byers, and J. Considine. Approximately uniform random sampling in sensor networks. In *Proceedings of the 1st International Workshop on Data Management for Sensor Networks (Held in Conjunction with VLDB 2004)*, pages 32–39, 2004.
- [6] F. Battiston, G. Cencetti, I. Iacopini, V. Latora, M. Lucas, A. Patania, J.-G. Young, and G. Petri. Networks beyond pairwise interactions: Structure and dynamics. *Physics Reports*, 874:1–92, 2020.
- [7] A. R. Benson. Three Hypergraph Eigenvector Centralities. *SIAM Journal on Mathematics of Data Science*, 1(2):293–312, 2019.
- [8] A. R. Benson, R. Abebe, M. T. Schaub, A. Jadbabaie, and J. Kleinberg. Simplicial closure and higher-order link prediction. *Proceedings of the National Academy of Sciences*, 115(48):E11221–E11230, Nov. 2018.
- [9] J.-C. Bermond, C. Delorme, and J.-J. Quisquater. Strategies for interconnection networks: Some methods from graph theory. *Journal of Parallel and Distributed Computing*, 3(4):433–449, 1986.
- [10] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah. Julia: A fresh approach to numerical computing. *SIAM review*, 59(1):65–98, 2017.
- [11] M. A. Bhuiyan, M. Rahman, M. Rahman, and M. Al Hasan. GUISE: Uniform Sampling of Graphlets for Large Graph Analysis. In *Proceedings of the 12th IEEE International Conference on Data Mining, ICDM’12*, pages 91–100, 2012.
- [12] G. Casiraghi, V. Nanumyan, I. Scholtes, and F. Schweitzer. From Relational Data to Graphs: Inferring Significant Links Using Generalized Hypergeometric Ensembles. In *Social Informatics*, volume 10540, pages 111–120, 2017.
- [13] Y. M. Chee, T. Etzion, H. M. Kiah, A. Vardy, V. K. Vu, and E. Yaakobi. Constrained de bruijn codes: Properties, enumeration, constructions, and applications. *arXiv:2005.03102*, 2020.
- [14] F. Chiericetti, A. Dasgupta, R. Kumar, S. Lattanzi, and T. Sarlós. On sampling nodes in a network. In *Proceedings of the 25th International Conference on World Wide Web, WWW’16*, pages 471–481, 2016.
- [15] R. Chikhi, A. Limasset, S. Jackman, J. T. Simpson, and P. Medvedev. On the Representation of De Bruijn Graphs. *Journal of Computational Biology*, 22(5):336–352, May 2015.

- [16] P. S. Chodrow. Configuration models of random hypergraphs. *Journal of Complex Networks*, 8(3), 2020.
- [17] C. Cooper, T. Radzik, and Y. Siantos. Estimating network parameters using random walks. *Social Network Analysis and Mining*, 4(1):168, Dec. 2014.
- [18] C. Cooper, T. Radzik, and Y. Siantos. Fast low-cost estimation of network properties using random walks. *Internet Mathematics*, 12(4):221–238, 2016.
- [19] L. d. F. Costa and G. Travieso. Exploring complex networks through random walks. *Physical Review E: Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics*, 75(1):016102, Jan. 2007.
- [20] J. Dean and S. Ghemawat. MapReduce: Simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113, Jan. 2008.
- [21] P. Erdős and A. Rényi. On the evolution of random graphs. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, 5(1):17–60, 1960.
- [22] T. S. Evans and R. Lambiotte. Line graphs, link partitions, and overlapping communities. *Physical Review E*, 80(1):016105, July 2009.
- [23] J. Fairbanks, M. Besançon, S. Simon, J. Hoffman, N. Eubank, and S. Karpinski. Juliagraphs/graphs.jl: an optimized graphs package for the julia programming language, 2021.
- [24] P. Faizian, M. A. Mollah, X. Yuan, Z. Alzaid, S. Pakin, and M. Lang. Random Regular Graph and Generalized De Bruijn Graph with k -Shortest Path Routing. *IEEE Transactions on Parallel and Distributed Systems*, 29(1):144–155, Jan. 2018.
- [25] K. V. Garimella, Z. Iqbal, M. A. Krause, S. Campino, M. Kekre, E. Drury, D. Kwiatkowski, J. M. Sá, T. E. Wellems, and G. McVean. Detection of simple and complex de novo mutations with multiple reference sequences. *Genome Research*, page genome;gr.255505.119v1, Aug. 2020.
- [26] E. N. Gilbert. Random graphs. *The Annals of Mathematical Statistics*, 30(4):1141–1144, 1959.
- [27] C. Gkantsidis, M. Mihail, and A. Saberi. Random walks in peer-to-peer networks: Algorithms and evaluation. *Performance Evaluation*, 63(3):241–263, 2006.
- [28] C. Gote, G. Casiraghi, F. Schweitzer, and I. Scholtes. Predicting Sequences of Traversed Nodes in Graphs using Network Models with Multiple Higher Orders. *arxiv:2007.06662*, July 2020.

- [29] H. Hartle, B. Klein, S. McCabe, A. Daniels, G. St-Onge, C. Murphy, and L. Hébert-Dufresne. Network Comparison and the Within-Ensemble Graph Distance. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 476(2243), Nov. 2020.
- [30] I. Iacopini, G. Petri, A. Barrat, and V. Latora. Simplicial models of social contagion. *Nature Communications*, 10(1):2485, Dec. 2019.
- [31] Z. Iqbal, M. Caccamo, I. Turner, P. Flicek, and G. McVean. De novo assembly and genotyping of variants using colored de Bruijn graphs. *Nature Genetics*, 44(2):226–232, 2012.
- [32] A. Jazayeri and C. C. Yang. Motif discovery algorithms in static and temporal networks: A survey. *Journal of Complex Networks*, 8(4), Dec. 2020.
- [33] D. Jurgens and T.-C. Lu. Temporal Motifs Reveal the Dynamics of Editor Interactions in Wikipedia. In *Proceedings of 2012 the International AAAI Conference on Web and Social Media*, ICWSM’12, 2012.
- [34] L. Kovanen, M. Karsai, K. Kaski, J. Kertész, and J. Saramäki. Temporal motifs in time-dependent networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2011(11):P11005, Nov. 2011.
- [35] L. Kovanen, K. Kaski, J. Kertész, and J. Saramäki. Temporal motifs reveal homophily, gender-specific patterns, and group talk in call sequences. *Proceedings of the National Academy of Sciences*, 110(45):18070–18075, 2013.
- [36] R. Lambiotte, M. Rosvall, and I. Scholtes. From networks to optimal higher-order models of complex systems. *Nature Physics*, 15(4):313–320, 2019.
- [37] T. LaRock. DeBruijnNets.jl software package. <https://www.github.com/tlarock/DeBruijnNets.jl>, 2021.
- [38] T. LaRock, V. Nanumyan, I. Scholtes, G. Casiraghi, T. Eliassi-Rad, and F. Schweitzer. HYPA: Efficient Detection of Path Anomalies in Time Series Data on Networks. In *Proceedings of the 2020 SIAM International Conference on Data Mining*, SDM’20, pages 460–468, 2020.
- [39] G. Lee, J. Ko, and K. Shin. Hypergraph Motifs: Concepts, Algorithms, and Discoveries. *Proceedings of the VLDB Endowment*, 13(12):2256–2269, Aug. 2020.
- [40] A. Lempel. On a Homomorphism of the de Bruijn Graph and its Applications to the Design of Feedback Shift Registers. *IEEE Transactions on Computers*, C-19(12):1204–1209, 1970.
- [41] A. Lempel. On extremal factors of the de Bruijn graph. *Journal of Combinatorial Theory, Series B*, 11(1):17–27, Aug. 1971.

- [42] P. Liu, V. Guarrasi, and A. E. Sariyüce. Temporal Network Motifs: Models, Limitations, Evaluation. *arxiv:2005.11817*, 2020.
- [43] D. Loguinov, J. Casas, and X. Wang. Graph-theoretic Analysis of Structured Peer-to-Peer Systems: Routing Distances and Fault Resilience. *IEEE/ACM Transactions on Networking*, 13(5):1107–1120, 2005.
- [44] X. Lu and S. Bressan. Sampling connected induced subgraphs uniformly at random. In *Proceeding of the 24th International Conference on Scientific and Statistical Database Management*, SSDBM’12, pages 195–212, 2012.
- [45] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon. Network Motifs: Simple Building Blocks of Complex Networks. *Science*, 298(5594):824–827, 2002.
- [46] J. Mykkeltveit. A proof of Golomb’s conjecture for the de Bruijn graph. *Journal of Combinatorial Theory, Series B*, 13(1):40–45, Aug. 1972.
- [47] A. Paranjape, A. R. Benson, and J. Leskovec. Motifs in Temporal Networks. In *Proceedings of the 10th ACM International Conference on Web Search and Data Mining*, WSDM’17, pages 601–610, 2017.
- [48] S. Patra and A. Mohapatra. Review of tools and algorithms for network motif discovery in biological networks. *IET Systems Biology*, 14(4):171–189, Aug. 2020.
- [49] G. Petri and A. Barrat. Simplicial Activity Driven Model. *Physical Review Letters*, 121(22):228301, Nov. 2018.
- [50] L. V. Petrovic and I. Scholtes. Counting Causal Paths in Big Times Series Data on Networks. *arxiv:905.11287*, 2019.
- [51] P. A. Pevzner, H. Tang, and M. S. Waterman. An Eulerian path approach to DNA fragment assembly. *Proceedings of the National Academy of Sciences*, 98(17):9748–9753, 2001.
- [52] G. E. Pibiri and R. Venturini. Handling Massive N-Gram Datasets Efficiently. *ACM Transactions on Information Systems*, 37(2):1–41, 2019.
- [53] B. Ribeiro and D. Towsley. Estimating and sampling graphs with multi-dimensional random walks. In *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*, IMC’10, pages 390–403, 2010.
- [54] B. Ribeiro, P. Wang, F. Murai, and D. Towsley. Sampling directed graphs with random walks. In *Proceedings of the 2012 IEEE INFOCOM*, pages 1692–1700, 2012.
- [55] P. Ribeiro, P. Paredes, M. E. P. Silva, D. Aparicio, and F. Silva. A Survey on Subgraph Counting: Concepts, Algorithms and Applications to Network Motifs and Graphlets. *arXiv:1910.13011*, Oct. 2019.

- [56] P. Ribeiro and F. Silva. G-Tries: A data structure for storing and finding subgraphs. *Data Mining and Knowledge Discovery*, 28(2):337–377, Mar. 2014.
- [57] R. A. Rossi, N. K. Ahmed, A. Carranza, D. Arbour, A. Rao, S. Kim, and E. Koh. Heterogeneous Network Motifs. *arxiv:1901.10026*, Jan. 2019.
- [58] J. Saramäki. Characterizing Motifs in Weighted Complex Networks. In *AIP Conference Proceedings*, volume 776, pages 108–117. AIP, 2005.
- [59] S. Sarkar, R. Guo, and P. Shakarian. Using network motifs to characterize temporal network evolution leading to diffusion inhibition. *Social Network Analysis and Mining*, 9(1):14:1–14:24, Dec. 2019.
- [60] I. Scholtes. When is a network a network?: Multi-order graphical model selection in pathways and temporal networks. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD’17, pages 1037–1046, 2017.
- [61] I. Scholtes, N. Wider, and A. Garas. Higher-order aggregate networks in the analysis of temporal networks: Path structures and centralities. *The European Physical Journal B*, 89(3), Mar. 2016.
- [62] I. Scholtes, N. Wider, R. Pfitzner, A. Garas, C. J. Tessone, and F. Schweitzer. Causality-driven slow-down and speed-up of diffusion in non-Markovian temporal networks. *Nature Communications*, 5(1), Dec. 2014.
- [63] A. Schwarze and M. A. Porter. Motifs for Processes on Networks. *SIAM Journal on Applied Dynamical Systems*, 20(4):2516–2557, Jan. 2021.
- [64] V. Sekara, A. Stopczynski, and S. Lehmann. Fundamental structures of dynamic social networks. *Proceedings of the National Academy of Sciences*, 113(36):9977–9982, 2016.
- [65] S. S. Shen-Orr, R. Milo, S. Mangan, and U. Alon. Network motifs in the Transcriptional regulation network of Escherichia coli. *Nature genetics*, 31(1):64, 2002.
- [66] R. Sinatra, D. Condorelli, and V. Latora. Networks of Motifs from Sequences of Symbols. *Physical Review Letters*, 105(17):178702, Oct. 2010.
- [67] S. Soundarajan, A. Tamersoy, E. B. Khalil, T. Eliassi-Rad, D. H. Chau, B. Gallagher, and K. Roundy. Generating Graph Snapshots from Streaming Edge Data. In *Proceedings of the 25th International World Wide Web Conference*, WWW’16, pages 109–110, 2016.
- [68] L. Stone, D. Simberloff, and Y. Artzy-Randrup. Network motifs and their origins. *PLOS Computational Biology*, 15(4):e1006749, Apr. 2019.

- [69] L. Torres, A. S. Blevins, D. Bassett, and T. Eliassi-Rad. The Why, How, and When of Representations for Complex Systems. *SIAM Review*, 63(3):435–485, 2021.
- [70] R. TransStat. Origin and destination survey database, 2018.
- [71] K. Tu, J. Li, D. Towsley, D. Braines, and L. D. Turner. Network Classification in Temporal Networks Using Motifs. *arXiv:1807.03733*, Aug. 2018.
- [72] J. Ugander, L. Backstrom, and J. Kleinberg. Subgraph Frequencies: Mapping the Empirical and Extremal Geography of Large Graph Collections. In *Proceedings of the 22nd International Conference on World Wide Web*, WWW’13, pages 1307–1318, 2013.
- [73] W. G. Underwood, A. Elliott, and M. Cucuringu. Motif-based spectral clustering of weighted directed networks. *Applied Network Science*, 5(1):62, Dec. 2020.
- [74] S. Wasserman and K. Faust. *Social Network Analysis: Methods and Applications*. Cambridge University Press, 1994.
- [75] R. West and J. Leskovec. Human wayfinding in information networks. In *Proceedings of the 21st International Conference on World Wide Web*, WWW’12, pages 619–628, 2012.
- [76] J. Xu, T. L. Wickramaratne, and N. V. Chawla. Representing higher-order dependencies in networks. *Science Advances*, 2(5):e1600028–e1600028, 2016.
- [77] Q. Xuan, H. Fang, C. Fu, and V. Filkov. Temporal motifs reveal collaboration patterns in online task-oriented networks. *Physical Review E: Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics*, 91(5):052813, May 2015.
- [78] X. Yan and J. Han. CloseGraph: Mining Closed Frequent Graph Patterns. In *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD’03, pages 286–295, 2003.
- [79] X. Yan, X. J. Zhou, and J. Han. Mining closed relational graphs with connectivity constraints. In *Proceeding of the 11th ACM SIGKDD International Conference on Knowledge Discovery in Data Mining*, KDD’05, pages 324–333, 2005.
- [80] S. Yoon, S. Lee, S.-H. Yook, and Y. Kim. Statistical properties of sampled networks by random walks. *Physical Review E: Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics*, 75(4):046114, Apr. 2007.
- [81] J.-G. Young, G. Petri, F. Vaccarino, and A. Patania. Construction of and efficient sampling from the simplicial configuration model. *Physical Review E*, 96(3):032312, 2017.

- [82] D. R. Zerbino and E. Birney. Velvet: Algorithms for de novo short read assembly using de Bruijn graphs. *Genome Research*, 18(5):821–829, Feb. 2008.

A Parallel Construction Algorithm

As mentioned in Section 3, the for loops in CAC can be parallelized. We note that constructing the DeBruijn graph from the list of sequences \mathcal{S} is the same as counting n-grams, a problem that has been extensively studied and parallelized [52]. In particular, we can do the counting within a Map-Reduce framework [20], where the mapping step splits each walk $w \in \mathcal{S}$ into length k segments (edges in the k th order DeBruijn graph), while simultaneously projecting each segment into the alphabet Σ . Then the Reduce step aggregates the count of each segment (i.e., edges of the DeBruijn Graph) as well as each projected motif. Algorithm 2 shows the pseudocode for this procedure.

Algorithm 2 Map and Reduce functions for simultaneously constructing a DeBruijn graph and counting motifs.

```

1: function MAP( $\mathcal{S} = (w_1, w_2, \dots, w_n)$ )
2:   for  $w_i$  in  $\mathcal{S}$  do
3:      $\ell = \text{length of } w_i$ 
4:      $\triangleright$  Split into length  $k$  segments and project each
5:     for  $j$  in  $0, \dots, \ell - k + 1$  do
6:        $\text{kseg} \leftarrow w_i[j, j + k]$ 
7:        $m \leftarrow \text{ProjectMotif}(\text{kseg})$ 
8:       Reduce( $\text{kseg}, 1$ )
9:       Reduce( $m, 1$ )
10:     $\triangleright$  Increment the count of  $x$ , which is either a sequence or a motif, by  $f$ 
11:  function REDUCE( $x, f$ )
12:    IncreaseCount( $x, f$ )

```

B Recovery of Network Structures

In this section we briefly examine the relationship between sampling uniformly from G_c^k and recovering network structure, both 1st and higher-order. In the first column of Figure 7, we show for each dataset the proportion of observed 1st-order nodes and edges recovered after sampling increasing number of walks from G_c^k . In both the flights (top) and Wikispeedia (bottom) datasets, all 1st-order nodes appear after a relatively small number of walks, while the whole edgeset is not sampled until about 75k walks have been sampled in the flights, and 500k walks sampled in Wikispeedia.

In the middle column of Figure 7, we show the proportion of observed and unobserved 3rd-order nodes observed as the number of walks sampled from G_c^k increases. Note that we analyze the 3rd-order rather than the 2nd-order network because unobserved nodes in the 2nd-order network would be unobserved edges in the 1st-order network, which will never appear by definition. After sampling 2 million walks uniformly at random, we discover all of the observed 3rd-order nodes in the flights dataset, but only about 60% of the unobserved nodes. In

the Wikispeedia data, after 2 million walks we have only sampled about 70% of the observed nodes and 50% of unobserved nodes.

Finally, in the last column of Figure 7, we show the same relationship as above but for 3rd-order edges. After 2 million samples, the proportion of sampled edges, both observed and unobserved, is less than 1% for both datasets. The relationship between walks sampled and edges observed is linear, which is expected since walks are being sampled uniformly, and the number of walks sampled is much smaller than the total number of possible walks. This means that on average each sample adds 1 to either the observed or unobserved count.

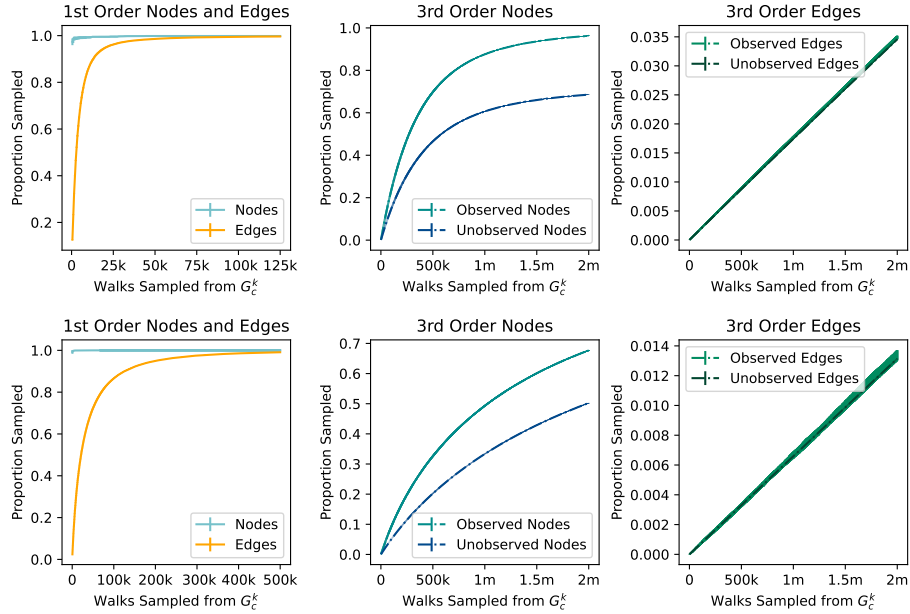


Figure 7: Relationships between the number of walks sampled uniformly from G_c^k and recovery of network structure in the flights (top) and Wikispeedia (bottom) datasets. Left panel: first-order structure is recovered relatively quickly, after about 75k walks in the flights and 500k walks in Wikispeedia. Middle panel: the proportion of observed 3rd-order nodes appearing in sampled walks reaches 100% in the flights dataset and around 70% in the Wikispeedia after 2 million samples, while 60% and 50% of unobserved nodes appear in the same samples. Right panel: the number of observed and unobserved 3rd-order edges increases linearly with the number of walks sampled.

C Approximation Procedures

We reduce the computational burden of sampling from G_c^k by using the following procedure to sample edges uniformly without generating the entire DeBruijn

graph. First, we sample a node u from the set of all nodes in V that are the source of at least one k -edge walk. Formally, we compute the number of walks originating from a node u by taking the row summation of the k th power of the adjacency matrix of the directed graph G , which we will denote as $w_u = \sum_v \mathbf{A}_{uv}^k$. We then sample from the set of nodes where $w_u > 0$ with probability proportional to w_u , and generate all k -edge walks starting from u by constructing all walks of length one—e.g., all of the directed edges (u, v) pointing away from u —and iteratively expanding the set of walks by one edge—e.g., appending each neighbor of v to the end of the previous walk—until the walks become longer than k -edges. Finally, we sample one of these walks uniformly at random.

We argue that the above sampling procedure is equivalent to sampling from all edges in G_c^k . We note that the probability of sampling a given first-order walk from G_c^k is simply the probability of choosing any of its m edges: $p_{e_i} = \frac{1}{m}$. In contrast, our procedure requires two steps. First, we sample a starting node u with probability $\frac{w_u}{\sum_{v \in V} w_v}$. Second, we sample one of the w_u k -edge walks starting from u uniformly. The probability of sampling an arbitrary walk w starting from any node u is the product of these two terms:

$$p_w = \frac{w_u}{\sum_{v \in V} w_v} \frac{1}{w_u} = \frac{1}{\sum_{v \in V} w_v} = \frac{1}{m} = p_{e_i}.$$

Thus the probability p_w of sampling a k -edge walk w using our two-step procedure is equal to the probability of sampling the same walk as an edge e_i from the complete DeBruijn graph G_c^k .

While the above procedure is more computationally tractable than constructing G_c^k at the outset, if the number of walks we want to compute is substantially larger than N , we will implicitly construct G_c^k anyway, since we will eventually compute all k -edge walks starting from all nodes u . To further reduce the computational burden, we use random walk simulations to sample. Instead of constructing *all* w_u k -edge walks starting from a node u , we simulate some number $x < w_u$ random walks, rejecting any walks with fewer than k edges (e.g., if a walker arrives at a node with no out-degree). Then we sample a walk from this subset of random walks. If a node is sampled again, we compute another set of x random walks, and union them with the previously computed set. Note that this procedure is no longer a uniform sample from the edges of G_c^k , since the probability of a random walk is determined by the out-degrees of its constituent nodes, meaning we are sampling a walk uniformly from a non-uniform sample of walks starting from u . We address this by sampling a random walk with probability proportional to the product of the out-degrees of the first k nodes of the walk. This means that walks, which are individually less likely because they visit multiple hubs with many choices, are weighted more highly by our sampling procedure. This procedure will also eventually construct G_c^k in its entirety, but we can now trade between speed and uniformity of the sample by setting x lower (faster) or higher (more accurate) compared to the true number of walks per node.

An even simpler approximation is to simulate M k -edge random walks on the first-order graph G , both with and without edge weights. This corresponds to comparing the observed DeBruijn graph to a version where all of the unobserved edges (including between nodes that were not observed in G^k) have weight proportional to their probability to be observed in G . This method also benefits from the familiarity of interpretation of random walks on graphs.

C.1 Numerical Convergence in Synthetic Data

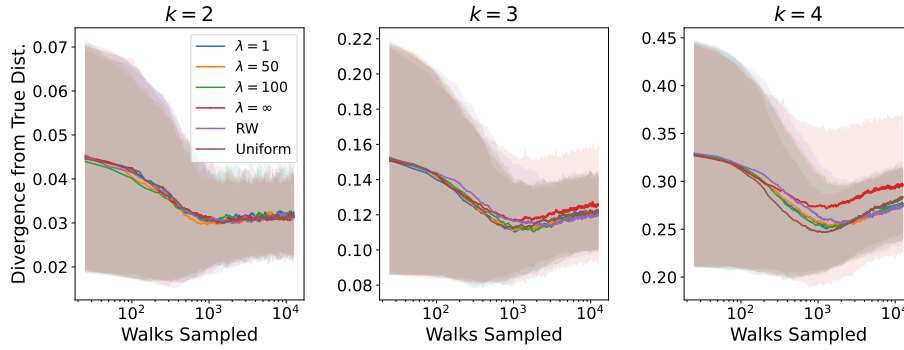


Figure 8: Convergence of sampled motif distribution towards the true motif distribution for orders $k = 2, 3, 4$. The horizontal axis shows the cumulative number of walks sampled. The vertical axis shows the KL-divergence between the true motif distribution and the sampled motif distribution after accumulating the corresponding number of walks on the horizontal axis. The parameter λ indicates the number of random walks computed per walk sampled for the approximate methods, with $\lambda = \infty$ corresponding to sampling one walk from all possible walks for each sampled node. “RW” means sampling a single random walk in the 1st-order graph G at every step. “Uniform” means sampling a walk uniformly from all possible k -edge walks based on G_c^k . All methods converge to the uniform distribution at similar rates, but as k increases, we see that the true uniform sampling method converges more quickly on average.

We evaluate the efficacy of the sampling procedures described above in the following way. First, we generate a random network G from $G_{N,p}$ [26, 21], with $N = 500$ and $p = 0.005$. For orders $k = 2, 3, 4$, we compute the true motif distribution by computing the complete DeBruijn graph of G and mapping all of its edges—all k -edge walks—to motifs, then normalize the motif counts to sum to 1. We then choose integers w and i , where i represents the number of iterations and w represents the number of walks sampled per iteration. For every iteration we sample w walks using each sampling method, map the walks to motifs, and accumulate the count of each motif for each method. Then we use the accumulated motif counts to compute a sample distribution and compute the KL-divergence between the sampled distribution up to iteration j and the

true distribution.

We show the results of this simulation using $w = 25$ and $i = 500$ in Figure 8, where each panel corresponds to an order k . In each plot of Figure 8, the horizontal axis represents the cumulative number of walks sampled so far—i.e., $w, w \cdot 2, \dots, w \cdot i$ —and the vertical axis represents the KL-divergence between the true motif distribution and the sampled distribution so far, averaged over 1000 samplings. Each line corresponds to a different sampling procedure: $\lambda = \infty$ refers to computing all k -edge walks for every sampled start node; $\lambda = 1, 10, 50, 100$ refer to computing λ walks per sampled node; “RW” refers to sampling a single random walk from the 1st-order graph G ; and “Uniform” refers to sampling one walk from all possible k -edge walks based on G_c^k . At $k = 2$, where there are only 2 sequential motifs, the KL-divergence is already close to the true distribution after only a small number of walks are sampled. However, for $k = 4$, we see that the uniform sampling has a slightly lower average for the first 1000 samples.

D Additional Results: Wikispeedia Unfinished and Flights Q2

In Figure 9, we reproduce the Wikispeedia - Finished results from Figure 5 on the left, and show the results for Wikispeedia - Unfinished on the right. The results are very similar overall, however there is 1 key difference worth highlighting. The directed triangle motif (6th row of Figure 9 appears at similar frequencies in both finished and unfinished games (around 500 observations), but in the finished games, the uniform ensemble suggests that the motif is either within expectation or slightly under-represented, while in the unfinished games the motif appears over-represented. This could indicate that triangles are indicative of a bad game, because the player makes two steps, gets frustrated with their options, then returns to an earlier node.

In Figure 10, we compare flight motifs between Q1 (left, reproduced from Figure 5) and Q2 (right) of 2020. The overall number of flights dropped considerably, likely due to the COVID-19 pandemic. Specifically, the use of the 2-node chain motif (2nd row in Figure 10) decreased both in terms of observed trips and in the extent to which it is underrepresented, potentially indicating that passengers were opting for 1-way trips during the pandemic. Future work should examine these patterns while controlling for year-to-year seasonality.

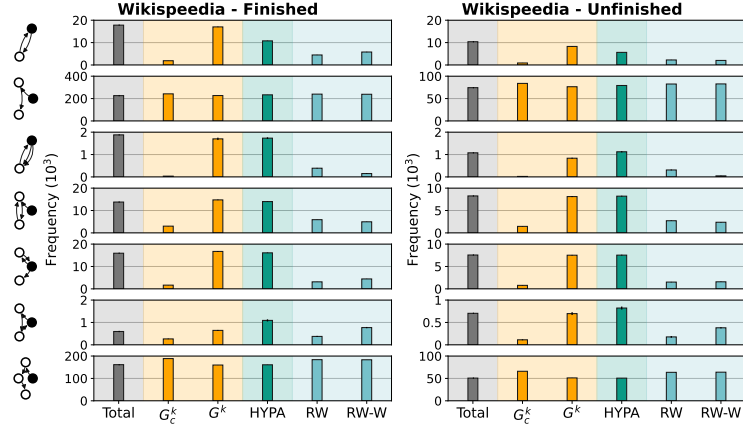


Figure 9: Sequential motif results for Finished Wikispeedia games (left, reproduced from Figure 5) and Unfinished games (right). Directed triangles (row 6 of the figure) play different roles in finished and unfinished games.

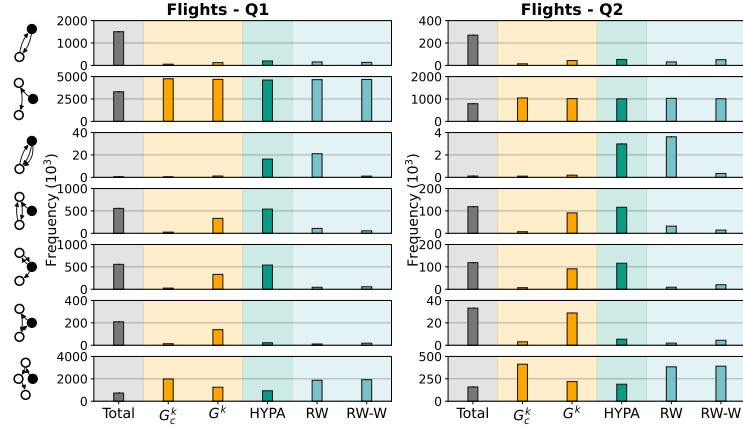


Figure 10: Sequential motif results for flights in Quarter 1 of 2020 (left, reproduced from Figure 5) compared to Quarter 2 (right). Comparing Q1 and Q2, we observe chain motifs (row 2) in Q2 at rates much closer to the expectation the null models. This might indicate a preference on the part of passengers for 1-way trips to destinations where they planned to stay during the pandemic, rather than round trips that are typical according to Q1 data.