

Predicting Sequences of Traversed Nodes in Graphs using Network Models with Multiple Higher Orders

Christoph Gote¹, Giona Casiraghi¹, Frank Schweitzer¹, and Ingo Scholtes^{2,3}

¹Chair of Systems Design, ETH Zurich, Zurich, Switzerland

²Chair of Data Analytics, University of Wuppertal, Wuppertal, Germany

³Data Analytics Group, University of Zurich, Zurich, Switzerland

Abstract

We propose a novel sequence prediction method for sequential data capturing node traversals in graphs. Our method builds on a statistical modelling framework that combines multiple higher-order network models into a single multi-order model. We develop a technique to fit such multi-order models in empirical sequential data and to select the optimal maximum order. Our framework facilitates both next-element and full sequence prediction given a sequence-prefix of any length. We evaluate our model based on six empirical data sets containing sequences from website navigation as well as public transport systems. The results show that our method out-performs state-of-the-art algorithms for next-element prediction. We further demonstrate the accuracy of our method during out-of-sample sequence prediction and validate that our method can scale to data sets with millions of sequences.

The analysis of data on complex networks provides essential insights into the structure and dynamics of social, technical, and biological systems. Apart from data that capture the *topology* of networked systems, i.e., which elements are directly connected via links, we increasingly have access to time-resolved, sequential data on *paths or trajectories in networks*. Examples include clickstream data generated by users who follow hyperlinks in information networks, data on information cascades propagating along friendship relations in online social networks, or temporal data capturing the travel routes of passengers in a transportation network. We can represent such path data as a multi set S of node sequences $v_1 \rightarrow v_2 \rightarrow \dots \rightarrow v_l$ of variable-length l , in a graph $G = (V, E)$ under the constraint that nodes $v_i \in V$ and transitions $(v_i, v_{i+1}) \in E$ for $i = 1, \dots, l - 1$. That is, different from general categorical sequence data, paths are constrained sequences of nodes where consecutive nodes are connected by a link in an underlying graph (cf. fig. 1a). Such data provide two complementary dimensions of information on networked systems: (i) the network, which determines, e.g., which connections passengers can use or which hyperlinks users can click, and (ii) path data, which determine along which specific paths passengers actually travel or how users sequentially navigate through web pages.

Recent works have shown that data on paths in networked systems expose *higher-order dependencies* between nodes, i.e., correlations in the sequence of nodes traversed on paths that cannot be explained by the network topology [5, 27, 43]. The modelling of patterns in such data is both a topical and a challenging problem. On the one hand, data on paths in networks rule out state-of-the-art sequential pattern mining algorithms that do not account for (i) the constraint that an underlying network limits the sequences of traversed nodes, and (ii) the fact that path data typically consists of many short node sequences with varying lengths. On the other hand, path data invalidates network analysis and graph mining techniques that neglect sequential patterns in the sequence of traversed nodes.

Addressing this challenge, recent works in data science and network analysis have proposed to use *higher-order models for paths in networks*, which generalise graph representations to higher-dimensional models

that capture sequential patterns in paths. Such models capture the *topology of paths* rather than the *topology of links*, e.g., by representing a path $v_1 \rightarrow v_2 \rightarrow v_3$ as a link between two *second-order nodes* (v_1, v_2) and (v_2, v_3) in a second-order network model. This representation generalises the concept of a link in standard graph representations, which can be viewed as a path of length two, to k -th order networks where links represent paths of length $k + 1$. Such generalisations of graph models to higher orders improve our ability to cluster and rank nodes [35, 37, 47], detect anomalies [28], and open new perspectives for network embedding [6, 36]. However, the modelling assumptions underlying those works hinder an application to the supervised prediction of paths, which is the motivation for our work.

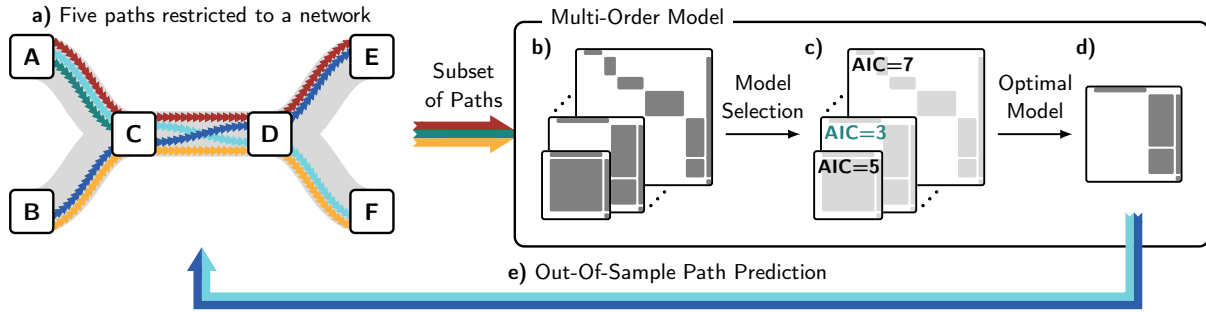


Figure 1: Overview of the proposed framework to predict paths in networks: Using paths observed in a network (a), MOGen learns a multi-order generative model (b) that accounts for (i) sequential patterns in the sequence of traversed nodes, (ii) endpoints of paths, and (iii) length distributions of paths in a training set. An information-based model selection algorithm (c) yields an optimal generative model (d) that can be used for an out-of-sample prediction of paths in a validation set (e).

Contributions: To address this issue, we introduce MOGen, a method to learn *generative multi-order models* in large collections of variable-length paths in networks. Our contributions are as follows. (i) Extending recent research on higher-, variable-, and multi-order models of paths, our framework combines Markov chains of multiple orders into a single generative model that accounts for (a) patterns in the sequence of traversed nodes, and (b) the constraint that sequences of traversed nodes are limited to an underlying network. (ii) Exposing a—to the best of our knowledge—previously unknown relation between higher-order network models and matrix representations of multi-layer networks [24], we show how our models can be formalised in terms of block adjacency and transition matrices. (iii) We propose an information-based model selection algorithm that detects the correct maximum model order needed to model sequential patterns in a given set of paths. We evaluate this algorithm in empirical data and show that it yields generalisable models that neither under- nor overfit sequential patterns. (iv) We develop a supervised learning technique to predict variable-length paths in networks. Different from existing algorithms, our method supports both predicting the next node based on a prefix of varying length (next-element prediction) as well as the out-of-sample prediction of full paths based on a training set. We apply our method to six data sets on human clickstreams and passenger itineraries in transportation networks and demonstrate its superior performance through a comparison to state-of-the-art sequence prediction algorithms. An overview of our method is shown in fig. 1.

The remainder of this article is structured as follows. First, we summarise the state-of-the-art the modelling of categorical sequences and paths and identify the research gap that is addressed by our work.

Subsequently we introduce our method. We then experimentally validate our method in six empirical data sets. We apply our method to the supervised prediction of paths, compare its performance to state-of-the-art sequence prediction algorithms, and evaluate our model selection algorithm. Finally, we summarise our conclusions.

1 Related Work

The modelling of categorical sequences has been addressed in the data mining and machine learning community. While we refer to [20, 42] for an exhaustive overview on sequential pattern mining, here we limit our discussion of related work to (i) methods to model and predict categorical sequences, and (ii) works that explicitly address the modelling of paths in a network.

Sequence Modelling and Prediction The modelling and prediction of categorical sequence data have important applications, e.g., in text mining, machine translation, DNA sequence analysis, clickstream analysis and human behavioural modelling. A large body of works has applied probabilistic graphical models like, e.g., Petri nets, hidden Markov models, decision trees, as well as higher- or variable-order Markov chains [7, 15, 40, 46] to model patterns in categorical sequences [20, 42]. Such sequence modelling techniques facilitate *next-element prediction*, which seeks to predict the next element(s) in a given sequence based on a model learned in training sequences. This problem occurs in a number of applications, e.g., recommender systems [21], speculative Web caching [9], or purchase prediction in eCommerce [25]. Such applications have been addressed with (higher-order) Markov chains that predict the next element in a sequence based on the previous k elements [8, 16, 17, 23, 33]. Other methods use Petri nets [3, 45], decision trees [13, 22, 29], graph and network models [26, 32], spectral learning [4], as well as neural networks and automata [42]. An advantage of those methods is that they can be applied to general data on categorical sequences. However, data on paths in networks have characteristics that violate the assumptions of some of those techniques. First, paths are *constrained* categorical sequences insofar as an underlying network restricts the possible sequences of traversed nodes. In terms of sequential pattern mining, this has the consequence that some frequently traversed node sequences are trivially due to the topology of the underlying network, while others can not be explained by the network alone. Second, different from data that contain a small number of long categorical sequences, path data typically confront us with a large unordered collection of *short and variable-length sequences*, where each sequence is one independent observation of, e.g., a user clickstream in the Web or an itinerary of a passenger in a transportation network. The structure of such data complicates the application of higher-order Markov chains with large orders k , which treat the first k nodes in each path as a prefix for the following transitions.

Higher-Order Models of Paths in Networks Several recent works highlight the need for *higher-order network models* to account for the special characteristics of sequential data on paths. [27] summarise the challenges of this research at the intersection of network analysis, data mining, and machine learning. The key idea behind those works is the use of higher-dimensional graph representations. The resulting higher-order network models resemble the state space extension that is the basis of higher-order Markov chains or high-dimensional De Bruijn graphs [12]. Building on this idea, [35] use a generalisation of

the flow compression algorithm InfoMap [34] to higher-order network models to detect temporal cluster structures in sequential data on networks. [39] show that the spectral properties of higher-order graph Laplacians can be used to improve the prediction of diffusion processes based on time-stamped and sequential data on networks. A number of works have generalised node centrality measures and ranking algorithms, demonstrating that higher-order models of paths improve our ability to identify important nodes in networks [35, 37, 38, 47]. Recent works used such models to capture event sequences in time-stamped network data [31], detect anomalies in data on paths in networks [28], or propose new network embedding and representation learning techniques [6, 36].

Those works generated insights into the topology of networked systems that cannot be obtained with standard network analysis and graph mining techniques. However, the modelling assumptions underlying those works also limit their application in many important settings. First, existing works have focused on higher-order models for patterns in the sequence of traversed nodes, rather than *generative models* that additionally model the endpoints of paths [27]. This limits the ability to predict variable-length paths, an important supervised learning problem that is the key motivation for our work. A second limitation is the choice of model selection techniques to learn *generalisable higher-order models* that provide an optimal balance between model complexity and explanatory power. Existing works either used heuristic techniques [35, 47] or analytical methods that rely on additional model characteristics like, e.g., a nested structure [37] that do not apply to generative models. Moreover, most existing works used models with a single order k , which limits their application in paths that exhibit correlations at multiple scales at once. While both model selection and sequence prediction have been addressed in the modelling of categorical sequences, we note that—due to the special characteristics outlined above—those methods cannot be immediately applied to path data.

In the remainder of this article, we develop a generative multi-order modelling framework for paths in networks that closes this gap. Different from state-of-the-art sequence modelling and prediction algorithms, our method explicitly addresses the modelling of unordered collections of variable-length node sequences constrained by a network. Different from previous works on higher- and multi-order models [35, 37, 47], we focus on a generative model that accounts for the endpoints of paths and is thus suitable for prediction tasks. By combining models of multiple higher-order into a multi-layered model structure, our method captures sequential patterns at multiple length scales simultaneously. We further develop an information-based model selection algorithm that allows detecting the optimal maximum order of the model without requiring the nestedness property that is the basis for the method proposed in [37].

2 Multi-Order Generative Model for Paths

We first formally define path data and present the generative modelling framework and model selection algorithm that are the main contributions of our work. We assume that we are given a multi-set of m paths $S = \{t_1, \dots, t_m\}$ on a network $G = (V, E)$ with nodes V and (directed) links $E \subseteq V \times V$. Each path $t_i = v_1 \rightarrow v_2 \rightarrow \dots \rightarrow v_{l_i}$ is an ordered sequence of traversed nodes under the constraint that $(v_j, v_{j+1}) \in E$ for $j \in [1, l_i - 1]$. We assume that paths in S have varying lengths l_i and we denote the maximum path length as l_{max} . Using a simple network-based model for the sequences of nodes traversed by paths, we could assume that paths are generated by a memoryless random walk that corresponds to a

first-order Markov chain. The nodes $v \in V$ of network G correspond to the state space of such a Markov chain, while links $(v, w) \in E$ in the network are associated with transition probabilities. This assumes that the next node v_{i+1} traversed by a path only depends on (i) the currently visited node v_i , and (ii) the links that determine which state transitions are possible. To capture higher-order dependencies in the sequence of traversed nodes that are not due to the network topology, we can generalise this model to a k -th order Markov chain, where the next node v_{i+k} depends on the k previously visited nodes v_{i-k+1}, \dots, v_i . We can view this model as a memoryless random walk in a k -th order network $G^k = (V^k, E^k)$, where (i) each k -th order node $(v_1, \dots, v_k) \in V^k$ is a sequence of k nodes where consecutive nodes are connected by links, and (ii) each k -th order link connects two higher-order nodes (v_1, \dots, v_k) and (w_1, \dots, w_k) such that $w_i = v_{i-1} \in V$. We note that a k -th order network can be constructed as the line graph of a $k - 1$ -th order network, a construction that adapts the construction of De Bruijn graphs to models for paths in a network [12, 37]. Using this modelling approach, a path $v_1 \rightarrow v_2 \rightarrow \dots \rightarrow v_l$ corresponds to $l - k + 1$ transitions along the edges of a k -th order network, which gives rise to the following sequence of $l - k + 1$ higher-order nodes:

$$(v_1, \dots, v_k) \rightarrow (v_2, \dots, v_{k+1}) \rightarrow \dots \rightarrow (v_{l-k+1}, \dots, v_l)$$

2.1 Multi-Order Models of Paths

Referring to the special characteristics of variable-length paths in S , the modelling approach above has a number of limitations. (i) Using a k -th order model, we can neither use the first k nodes in each path to fit the transition probabilities of the model, nor can we use the model to perform a next-element prediction for the first k nodes on a path. (ii) The fact that the number of observations used to fit the transition probabilities of a k -th order model differs across different orders k complicates model selection, i.e., the selection of the optimal order k to model a data set. (iii) The model neglects information on the start and end point of nodes, which hinders its application to an out-of-sample prediction of full paths, i.e. from start to end. (iv) Finally, a model with a single order k is limited to capture sequential patterns of a single length k , which is why recent works argued for variable- and multi-order modelling techniques [33, 37, 47].

To overcome those limitations we introduce a generative model that combines (i) transitions between higher-order models of multiple orders up to a maximum order K and (ii) special transitions that capture the start and end points of paths. For this, we introduce a special *initial state* $*$, with transitions $* \rightarrow v$ to first-order nodes $v \in V$ that mark the start of a path in node v . For all orders $1 \leq k \leq K$, we further add transitions from k -th order nodes $(v_1, \dots, v_k) \rightarrow \dagger$ to a special *terminal state* \dagger , which captures the termination of a path. With this extension, a path $v_1 \rightarrow v_2 \rightarrow \dots \rightarrow v_l$ corresponds to $l + 2$ transitions in a *multi-order network* with maximum order K , which gives rise to the following sequence of l higher-order nodes

$$* \rightarrow v_1 \rightarrow (v_1, v_2) \rightarrow (v_1, v_2, v_3) \rightarrow \dots \rightarrow (v_{l-K+1}, \dots, v_l) \rightarrow \dagger$$

with two additional transitions from/to the special states $*$ and \dagger to model the start and end of the path. Considering the coding of memory prefixes in the extended state space of higher-order Markov chains, we note that each transition in this multi-order model with maximum order K increases the memory prefix by one node, up to a maximum memory of length K . In addition to transitions *within* a K -th order model, we obtain a multi-order model that includes transitions between the nodes of a k -th order model

and the nodes of a $k + 1$ -th order model for all orders $k < K$. Transitions from/to initial state $*$ and terminal state \dagger explicitly model the start and end points of paths, which is the basis for an end-to-end, out-of-sample prediction of full paths.

We can mathematically represent such a multi-order network model with maximum order K by means of a *multi-order adjacency matrix* $\mathbf{A}^{(K)}$, whose overall structure is shown in fig. 2a. This representation is inspired by so-called *supra-adjacency matrices*, which have recently been proposed to represent the interconnected topologies of multi-layer networks [24]. A multi-order adjacency matrix consists of blocks $\mathbf{A}_{k-1,k}$, whose entries count the observed transitions between $k - 1$ -th and k -th order nodes in a given multi set S of paths. The entries in the top-left block $\mathbf{A}_{0,1} \in \mathbf{R}^{1 \times n}$ count transitions $(*, v)$ for all nodes $v \in V$, i.e., the distribution of start nodes for all paths. The entries of the bottom-right block \mathbf{A}^\dagger count *terminal transitions* $(v_1, \dots, v_k) \rightarrow \dagger$ for all orders $1 \leq k \leq K$. For paths with lengths $l_i > K$, the entries in the bottom matrix $\mathbf{A}_{K,K}$ count the transitions between the nodes in a K -th order network model. An example for a multi-order adjacency matrix that corresponds to a toy example where each path shown in fig. 1a is observed 10 times is shown in fig. 2a. We note that structural zeros in $\mathbf{A}^{(K)}$ (greyed out in fig. 2b) represent transitions that are *impossible* based on the underlying network topology, which are different from zero entries that represent possible but *unobserved* transitions (see zeros in highlighted blocks in fig. 2b). The only model parameter K determines the maximum order of the multi-order model, which resembles the maximum memory length in a Markov chain. We note that, under the assumption that the paths in a multi set $S = \{t_1, \dots, t_m\}$ with maximum path length l_{max} are generated independently of each other, the multi-order adjacency matrix $\mathbf{A}^{(K)}$ is a lossless representation of S iff $K \geq l_{max}$.

To facilitate prediction tasks, we use the multi-order adjacency matrix to define a *probabilistic generative model* of variable-length paths. This model is fully described by a *multi-order transition matrix* $\mathbf{T}^{(K)}$, which we obtain by a row normalisation of $\mathbf{A}^{(K)}$. This yields a stochastic matrix $\mathbf{T}^{(K)}$, whose entries give the transition probabilities in a probabilistic model for paths. The transition matrix of a generative multi-order model with maximum order $K = 3$ that corresponds to the toy example in fig. 1a is shown in fig. 2c. We are generally interested in maximally parsimonious models with maximum order $K \ll l_{max}$ that still capture relevant patterns in the sequence of nodes traversed by paths. Different from the lossless representation $\mathbf{A}^{(K)}$ obtained for $K > l_{max}$, we can view such a model as a lossy compression or summarisation of paths, based on the sequential patterns in S . In the following, we show how we can adapt information-theoretic model selection to choose the optimal maximum order K of a multi-order model that (i) yields an optimal summarisation of paths, and (ii) can be used to predict paths.

2.2 Model Selection

To choose a multi-order model with an optimal maximum order K for a given set of paths S , we must first evaluate how well a given model describes a data set. Taking a statistical inference perspective, we first define the probability of a single observed path $t = v_1 \rightarrow \dots \rightarrow v_l$ given a multi-order model described

A likelihood function for a multi-order model given a multi set S of m independent paths can then be obtained by taking the product over all path probabilities, i.e.,

$$\mathcal{L}(\mathbf{T}^{(K)}|S) = \prod_{t \in S} P(t|\mathbf{T}^{(K)}) \quad (1)$$

We highlight that the specific definition of the multi-order model allows the comparison of likelihoods across different maximum orders K . This is due to the fact that the number of observed transitions that enter the likelihood calculation in eq. (1) is independent of the maximum order K , which is not the case for a standard higher-order Markov chain model. We further note that, due to the inclusion of the special transitions to/from start and terminal states, the resulting models are not nested, i.e., a multi-order model with maximum order K is not contained in the parameter space of a more complex model with maximum order $K' > K$. While this difference to previous works on multi-order models is crucial to apply the model to the path prediction problem, it rules out the application of the likelihood-ratio based model selection proposed in [37].

We address this problem with an information-theoretic approach to model selection. We specifically use Akaike’s Information Criterion (AIC) [1] to balance the explanatory power of a multi-order model for a data set S with the complexity of the model, which is measured in terms of the degrees of freedoms in the parameter space. The AIC value can be interpreted as an estimate for the expected relative Kullback-Leibner (KL) divergence between the fitted model and the unknown true process that generates the observed data. This corresponds to the expected loss of information in bits if we were to replace the unknown process that generates the data by our model. The best model among a set of candidate models is thus the one that minimises Akaike’s Information Criterion. For a multi-order model, we can calculate the AIC as

$$\text{AIC}(\mathbf{T}^{(K)}) = 2d(\mathbf{T}^{(K)}) - 2 \ln \left\{ \mathcal{L}(\mathbf{T}^{(K)}|S) \right\} \quad (2)$$

where $d(\mathbf{T}^{(K)})$ are the degrees of freedom of the model defined by $\mathbf{T}^{(K)}$. Different from existing works that have used the AIC to determine the optimal order of higher-order Markov chain models of categorical sequence data [40, 41], we must account for (i) structural zeros in the transition matrix $T^{(K)}$ that result from the fact that we model paths restricted to a network, (ii) the multi-layer structure of our models that combine multiple orders k , and (iii) the additional model parameters used to model the transitions from/to the initial and terminal states. This implies that the degrees of freedom of a multi-order model of paths depend on the topology of the underlying network, which we can encode in terms of a binary adjacency matrix \mathcal{A} .¹ The possible paths of exactly length k in such a network are encoded by the non-zero entries in the k -th power of the binary adjacency matrix, which allows us to calculate the degrees of freedom of a model with maximum order K as:

$$d(\mathbf{T}^{(K)}) = \sum_{k=1}^K \sum_{i,j} (\mathcal{A}^k)_{ij} + |V| - 1 \quad (3)$$

¹If we assume that all possible transitions along the links of the underlying network have been observed at least once, we can directly derive \mathcal{A} from the training data.

Table 1: Summary statistics for the six data sets used to validate MOGen.

| | paths | | nodes on path | | | | network topology | | |
|---|-----------|--------|---------------|--------|-----|-----|------------------|--------|---------|
| | total | unique | mean | median | min | max | nodes | links | density |
| BMS1 : online retailer clickstreams [11] | 59'601 | 18'473 | 2.51 | 1 | 1 | 267 | 497 | 15'387 | 6.24% |
| FIFA : FIFA World Cup 98 server logs [2] | 20'450 | 20'053 | 36.24 | 31 | 9 | 100 | 2'990 | 73'530 | 0.82% |
| MSNBC : news website clickstreams [14] | 31'790 | 30'247 | 13.33 | 11 | 9 | 100 | 17 | 270 | 93.43% |
| WIKI : Wikispeedia clickstreams [46] | 76'193 | 68'784 | 6.25 | 5 | 1 | 435 | 4'179 | 69'800 | 0.40% |
| AIR : US flight itineraries [44] | 286'810 | 60'228 | 4.19 | 5 | 2 | 14 | 175 | 1'598 | 5.24% |
| TUBE : London Tube itineraries [30] | 4'295'731 | 32'313 | 7.86 | 7 | 2 | 36 | 276 | 663 | 0.87% |

Ignoring constant terms that do not depend on the maximum order K , using eq. (2) we can estimate the optimal maximum order \hat{K} for a data set S as follows:

$$\hat{K} := \arg \min_K \sum_{k=1}^K \sum_{i,j} (\mathcal{A}^k)_{ij} - \ln \left\{ \mathcal{L}(\mathbf{T}^{(K)} | S) \right\}. \quad (4)$$

To illustrate our method in the running toy example shown in fig. 1a, here we present results for a simple experiment, in which we vary the observed frequencies of the five (colour-coded) paths shown in fig. 1a. We apply the model selection to different observed frequencies of paths and report the estimated optimal parameter \hat{K} . Due to the topology of the network shown in fig. 1a, paths that start in node A or B and continue to C and D either terminate in node E or node F . In order to model paths whose probability to terminate in E or F depends on whether they start in A or B , a multi-order model with maximum order three is needed. Such a dependency is introduced if the paths $A \rightarrow C \rightarrow D \rightarrow E$ and $B \rightarrow C \rightarrow D \rightarrow F$ occur more or less frequently than the paths $B \rightarrow C \rightarrow D \rightarrow E$ and $A \rightarrow C \rightarrow D \rightarrow F$. Figure 2d shows that in this case our model selection algorithm yields the correct optimal order of three. If, however, the terminal node of a path is independent of the start node, a model with order two is sufficient in this toy example. This is the case if the frequencies of all four paths are similar, in which case our model selection algorithm yields the correct optimal order of two (cf. fig. 2d). This is illustrated by the identical values in $\mathbf{T}_{3,3}^{(3)}$ in fig. 2c. We note that those results merely serve to illustrate our method in the running toy example. In the coming section, we validate our method through a cross-validation experiment in empirical data sets. The results of this experiment show that the best prediction performance is obtained for the parameter \hat{K} that is estimated using the proposed model selection algorithm. This implies that, different from existing techniques, the parameter K can be directly estimated from the data, which effectively turns our method into a parameter-free approach.

3 Experimental Validation

We experimentally validate MOGen in six empirical data sets containing (i) user clickstreams on the Web, and (ii) travel itineraries of passengers in a train and an airline network: BMS1 contains 59'601 clickstreams of customers of the web retailer *Gazelle.com* [11]. FIFA captures 20'450 clickstreams of users of the '98 FIFA World Cup website [2]. MSNBC contains 31'790 clickstreams of page categories on the

MSNBC news website [14]. WIKI captures navigation paths of 76'193 users in the Wikipedia article graph, who were playing the game Wikispeedia [46]. AIR contains 286'810 flight itineraries of passengers travelling on routes between US airports in 2001 [44]. TUBE captures 4'295'731 itineraries of London Tube passengers [30]. The raw data for all data sets are freely available online. Summary statistics are provided in table 1. The data sets BMS1, FIFA and MSNBC were obtained from the repository published with [19]².

In the following sections, we validate our method as follows: Addressing an important application in sequence modelling, we first apply MOGen to predict the next node on a path, comparing its performance to six state-of-the-art methods. Highlighting a major advantage of our method over existing techniques, we then show that we can use the proposed model selection algorithm to compute the parameter K for which our model offers the best performance directly from the data. This avoids computationally expensive parameter search algorithms that have been used by competing sequence prediction methods. Showcasing another contribution of our work, we further evaluate our method in a scenario where we target an out-of-sample prediction of full paths. We finally report results that show the scalability of our method in real data.

3.1 Next-Element Prediction

We first evaluate the performance of our method in a next-element prediction setting. Given a prefix of k nodes (v_1, \dots, v_k) previously traversed by a path, we want to predict which node v_{k+1} in the network is visited next. We evaluate our method in a cross-validation experiment for the six empirical data sets outlined above. Leveraging the fact that our model captures terminal nodes on paths, we consider the special state \dagger as a prediction of a path termination. We include this in the evaluation, because the prediction of terminal nodes is crucial in data on variable-length path data, e.g., to predict where users exit a website, where passengers end their itinerary, or whether a purchase is made in an online shop. We compare the performance of MOGen to state-of-the-art sequence prediction algorithms that have addressed next-element prediction in categorical sequences. Our choice of baseline methods is informed by the review of [42], which compares the performance of state-of-the-art methods in several data sets. The authors find that AKOM [33] outperforms other algorithms in three out of four tested data sets. Conversely, [22] show that CPT+ outperforms AKOM in terms of the accuracy of the next element prediction. We thus use both CPT+ and AKOM as baseline, utilising the implementations provided in [19]. To enable a fair comparison for empty prefixes, we modified the implementation of AKOM to return a prediction based on the frequency of elements in the training data. This is described in [33], however the implementation in [19] returns an empty prediction instead. CPT+ predicts the next element in a sequence based on an internal assignment of weights to candidate elements, i.e., the prediction is not based on probabilities. We thus normalise the weights assigned by CPT+ to obtain a probabilistic prediction. While this allows us to compare the performance of CPT+ with the other (probabilistic) methods, we highlight that the evaluation in [22] is based on prediction accuracy. In addition to those sequence modelling techniques, we compare our method to a probabilistic prediction derived from (i) a random walk in a higher-order network with a single order k (NET), and (ii) the multi-order graphical model (MOG) introduced in [37]. For $k = 1$, NET is identical to a prediction generated by a random

²see <https://www.philippe-fournier-viger.com/spmf/index.php?link=datasets.php>

Table 2: Next-element prediction performance of all models.

| a) Cross-entropy loss [bit] for prefixes up to length 6 | | | | | | |
|---|-------------|-------------|-------------|-------------|-------------|-------------|
| | BMS1 | FIFA | MSNBC | WIKI | AIR | TUBE |
| MOGen | 6.22 | 7.49 | 2.72 | 7.83 | 4.20 | 1.81 |
| AKOM | 19.51 | 8.67 | 11.30 | 14.37 | 13.27 | 7.09 |
| CPT+ | 19.19 | 13.35 | 13.65 | 17.44 | 14.83 | 10.42 |
| NET | 19.51 | 8.84 | 11.44 | 14.36 | 13.34 | 7.58 |
| MOM | 19.51 | 8.66 | 11.29 | 14.36 | 13.20 | 6.35 |
| RND | 19.56 | 9.69 | 12.36 | 15.81 | 14.12 | 11.94 |
| b) MOGen: detected and best performing maximum order | | | | | | |
| | 1 | 1 | 2 | 1 | 2 | 6 |
| detected | 1 | 1 | 2 | 1 | 2 | 6 |
| best | 1 | 1 | 2 | 1 | 2 | 6 |

walk in the underlying network. Note that, despite a similarity in name, the method proposed in the present work considerably differs from MOM. First, MOM does not model the endpoint of paths, which limits its ability to predict variable-length paths. Second, it uses a model fitting approach that estimates transition probabilities based on sub-path frequencies rather than the transition probabilities used in our approach. Finally, it uses a model selection that builds on a nested model structure, which does not hold for our modelling framework. In addition to those network-based models, we use a naive (parameter-free) random baseline (RND) that simply predicts the next node based on the relative frequency of nodes in the training set, i.e., a prediction that neither uses a Markov chain nor the network topology. For AKOM, MOM, NET, and MOGen we report the results for the optimal (maximum) order parameter obtained via a grid search. For CPT+, we use the default parameters chosen in [22].

We split the paths in each data set into a training (90%) and validation (10%) set and evaluate the prediction performance of all models in terms of the *cross-entropy loss function*

$$H(p, q) := - \sum_{v \in V} p(v) \log q(v) \tag{5}$$

where, for a given prefix (v_1, \dots, v_k) , $q(v)$ denotes the probability of the next element (i.e., target node) v obtained from a model trained on the training set and $p(v)$ is the true distribution underlying the data. Since the true distribution is unknown for the empirical data set, we compute $p(v)$ from the distribution of actual next nodes in the validation set. This yields the log-loss function commonly used in the cross-validation of probabilistic multi-class prediction [10]. In the first evaluation, we further consider predictions for multiple prefix lengths up to a maximum length of six. We consider those as multiple samples and assign equal sample weights such that sample weights for any given target node with different prefix lengths sum to one. For the resulting loss function, a value of zero corresponds to a perfect prediction, which would allow us to replace the true distribution in the data by the prediction without information loss. Non-zero values quantify the information loss in bits, i.e., larger values correspond to worse prediction performance.

In table 2a we report the cross-entropy loss (in bits) for the six empirical data sets and the six prediction methods described above. For each algorithm, we report the performance for the parameters that yield

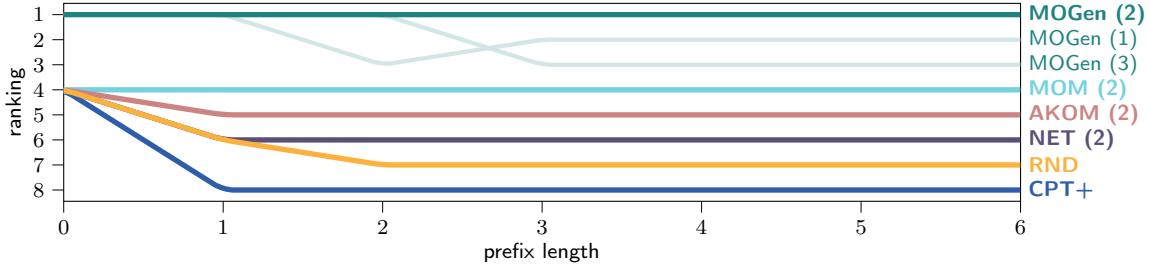


Figure 3: Prediction performance on different prefix lengths for MSNBC.

the smallest cross-entropy loss. The results show that MOGen outperforms all other methods for all of the six data sets. For most of the data sets, we further observe a considerable difference in the cross-entropy loss function. We hypothesise that this large difference in performance is due to the fact that different from other methods MOGen (i) explicitly models varying-length paths in a network, and (ii) is able to predict the termination of paths. This highlights the main contribution of our work, which is a model that specifically accounts for the characteristics of data on paths in networks. We further note that MON and NET yield identical results for BMS1 and WIKI. The reason for this is that in those two data sets both MOM and NET yield the best prediction performance for a (maximum) order of one, in which case a multi-order model is identical to a first-order model. We observe the smallest cross-entropy loss for MSNBC, AIR, TUBE, which we hypothesise is due to the —relative to the size of the underlying network— large number of observed paths. FIFA, WIKI, and BMS1 generally yield the largest cross-entropy loss across all methods, which is likely to be due to the large network topologies and a relatively small number of observations, which hinders a reliable detection of generalisable sequential patterns.

We finally observe that CPT+ shows the worst performance of all methods and it is even outperformed by the random baseline in all but one data sets. This is likely due to the fact that we evaluate the methods based on probabilities assigned to the next elements, which is not the originally intended use of CPT+. This is in line with the fact that a recent review found that the performance of CPT+ is inferior to the probabilistic method AKOM, while [22] have found the *accuracy* of predictions generated by CPT+ to be superior to AKOM.

Table 2 reports the prediction performance of methods across multiple prefix lengths and for the parametrisation that yields the best performance. A further interesting question is how the performance of methods compares for different prefix lengths and different parameters. To answer this question, in fig. 3 we show a ranking of methods for different prefix lengths used in the prediction in the MSBNC data set. For MOGen we show the performance of all considered parameters from $K = 1$ to $K = 3$ (MOGen(1) to MOGen(3)), while for other methods we show order parameter that yields the best performance (shown in brackets). The results show that MOGen outperforms other methods across all prefix lengths and for all parameters K .

3.2 Validation of Model Selection

All results reported in the previous section have been obtained by performing a grid-search of model parameters for which the models provide the best prediction performance. In the case of MOGen, the only parameter of our method is the maximum order K of the multi-order model. As explained previously, our method provides a model selection algorithm that enables us to estimate the optimal value of the maximum order \hat{K} directly from the data. This avoids the need for a computationally expensive grid search for the optimal parameter and effectively makes our method parameter-free, which is a major advantage over competing methods. In table 2b, we report the output of the AIC-based model selection of MOGen for the six empirical data set (top row), comparing it to the maximum order K that yields the best prediction performance (bottom row). For the clickstream data BMS1, FIFA, and WIKI we obtain an optimal maximum order of $K = 1$. For MSNBC, AIR, and TUBE we obtain optimal values of $K = 2$ and $K = 6$ respectively, which highlights the need for a higher-order model to capture sequential patterns. For all analysed data sets, the optimal order detected by the model selection provides the best prediction performance in a cross-validation, which confirms that our method yields models that neither underfit nor overfit sequential patterns in paths.

3.3 Out-Of-Sample Prediction of Paths

An important motivation of our method is its ability to perform an out-of-sample prediction of full paths based on a training set. Different from the next-element prediction addressed earlier, this addresses applications where we have access to a set of paths that allows us to learn a multi-order generative model. The task is then to predict a set of variable-length paths (from the start to the terminal nodes) that occur in a validation set. This problem has a number of interesting applications, e.g., speculative caching in web servers, the prediction of information cascades, or demand prediction in transportation systems based on past passenger itineraries. Addressing the latter scenario, we use MOGen to predict passenger itineraries in the London Tube based on a training set. We evaluate the performance of our method in terms of a binary classification experiment, where we use a training set to learn an optimal multi-order model and then use the learned model to predict the most frequently occurring paths in a validation set. The prediction of most frequent paths utilises the fact that our method learns a generative model that includes the start and end points of paths, i.e., we use our model to generate a new set S' of paths with variable lengths. We rank paths in the generated set S' based on their frequency and use the top N paths as prediction of the top 10% most frequent paths in the validation set. We interpret this prediction as a binary classification and repeat the prediction for different discrimination thresholds N . We then compute true positive/false positive rates for different values of N , which can be used to compute the area under a ROC (receiver operating characteristic) curve. The result of this out-of-sample prediction is shown in fig. 4. The ROC curve indicates a high accuracy of our method, which is substantiated by a resulting AUC score of 0.95.

3.4 Scalability

In a final experiment, we evaluate the scalability of our method. Our implementation parallelises the training of a multi-order model by splitting paths into multiple sets and simultaneously computing entries




Figure 4: Out-of-sample path prediction with MOGen. ROC and AUC for a 99/1 training/validation split for TUBE.

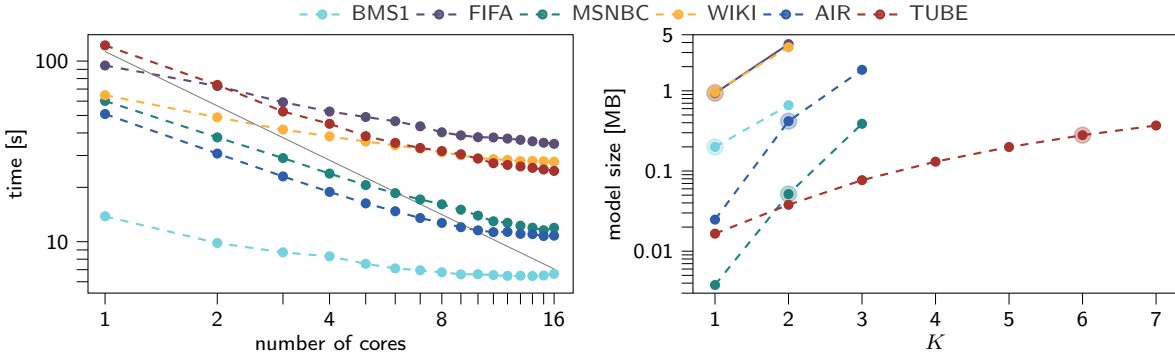


Figure 5: Mean time for training model selection, as well as model size for all data sets (5 iterations). The grey line in the left panel indicates a linear speedup. Highlighted models in the right panel are the ones picked through model selection.

of the multi-order adjacency matrix $\mathbf{A}^{(K)}$ and transition matrix $\mathbf{T}^{(K)}$ for each of the sets, and finally aggregating those entries. We speed up the computation of the model likelihood (cf. eq. (1)) by computing the probabilities of paths in parallel. The left panel of fig. 5 shows the time (in seconds) that is required to (i) train multi-order models for multiple parameters K , and (ii) select the optimal parameter \hat{K} using the method described in eq. (4) for different numbers of processing cores³. The resulting model sizes (in MegaBytes) for different orders K are shown in the right panel of fig. 5.

The results show that the time required to learn an optimal model primarily depends on the number

³Results were obtained using a 16-core Intel Core i9-7960X processor.

of unique paths as well as their length (cf. table 1). We also observe a near linear reduction of the training time as we increase the number of available processing cores (grey line refers to an optimal linear speedup). The deviation from the linear reduction is likely to originate from the non-parallel computation of the matrix powers of \mathcal{A} , which is required to compute the degrees of freedom. The size of the models generally increases super-linearly with the maximum order K . The scaling of model size depends on the size and algebraic connectivity [18] of the underlying topology, which determines the growth of the degrees of freedom with the parameter K . This is illustrated by the TUBE data set, where—due to a sparsely connected network topology—the increase in model size is lowest as the order increases. All models selected by MOGen are below 1 MB in size.

4 Conclusion

The growing availability of temporal data provides interesting opportunities for new data science techniques that address the modelling and prediction of paths in networks. We have shown that such data has unique characteristics that question the application of state-of-the-art sequential pattern mining, sequence modelling, and prediction algorithms. Moreover, data on paths contain information on the topology of interactions in networked systems that cannot be extracted with standard network analysis and graph mining techniques. Addressing this gap, we propose MOGen, a modelling and prediction framework that is based on multi-order generative models of paths observed in a network. Different from existing sequence mining and modelling frameworks, MOGen explicitly models the start- and endpoints of variable-length paths as well as patterns in node sequences that are restricted by an underlying network topology. The importance of those features is demonstrated by the fact that MOGen outperforms state-of-the-art sequence prediction techniques in terms of next-element prediction in six empirical data sets. Going beyond previous works on categorical sequence prediction, we further show that our method can be used for the out-of-sample prediction of full paths, e.g., variable-length paths from start to end, using passenger itineraries in a real transportation network. This has interesting applications, for example, in terms of demand prediction for smart mobility scenarios, that we currently explore in a collaboration with a transportation provider. A major advantage of MOGen compared to competing methods is that it provides an information-based model selection algorithm to determine the optimal order parameter K directly from data. An experimental evaluation in empirical data sets proves that MOGen is able to automatically learn models that generalise best to unseen data, which effectively turns it into a parameter-free method. Apart from those contributions to the foundation of sequential pattern mining in temporal network data, we provide an Open Source parallel implementation of our method⁴ and evaluate its scalability in real data sets.

Acknowledgements

The authors acknowledge discussions with Yan Zhang on an early version of this work. The author Ingo Scholtes acknowledges funding through the Swiss National Science Foundation grant 176938 as well as through the project bergisch.smart.mobility, funded by the German state of North Rhine-Westphalia.

⁴Integrated in <https://github.com/pathpy/pathpy>

References

- [1] Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. In: *2nd International Symposium on Information Theory*. pp. 267–281.
- [2] Arlitt, M.; Jin, T. (2000). A workload characterization study of the 1998 world cup web site. *IEEE network* **14(3)**, 30–37.
- [3] Augusto, A.; Conforti, R.; Dumas, M.; La Rosa, M. (2017). Split miner: Discovering accurate and simple business process models from event logs. In: *2017 IEEE International Conference on Data Mining (ICDM)*. IEEE, pp. 1–10.
- [4] Balle, B.; Carreras, X.; Luque, F. M.; Quattoni, A. (2014). Spectral learning of weighted automata. *Machine learning* **96(1-2)**, 33–63.
- [5] Battiston, F.; Cencetti, G.; Iacopini, I.; Latora, V.; Lucas, M.; Patania, A.; Young, J.-G.; Petri, G. (2020). Networks beyond pairwise interactions: structure and dynamics. *Physics Reports* .
- [6] Belth, C.; Kamran, F.; Tjandra, D.; Koutra, D. (2019). When to remember where you came from: node representation learning in higher-order networks. In: *Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. pp. 222–225.
- [7] Benson, A. R.; Gleich, D. F.; Lim, L.-H. (2017). The Spacey Random Walk: A Stochastic Process for Higher-Order Data. *SIAM Review* **59(2)**, 321–345.
- [8] Bernhard, S. D.; Leung, C. K.; Reimer, V. J.; Westlake, J. (2016). Clickstream prediction using sequential stream mining techniques with Markov chains. In: *Proceedings of the 20th International Database Engineering & Applications Symposium*. pp. 24–33.
- [9] Bestavros, A. (1995). Using speculation to reduce server load and service time on the WWW. In: *Proceedings of the fourth international conference on Information and knowledge management*. pp. 403–410.
- [10] Bishop, C. M. (2006). *Pattern recognition and machine learning*. springer.
- [11] Brodley, C.; Kohavi, R. (2000). KDD-Cup 2000 homepage.
- [12] Bruijn, de, N. (1946). A combinatorial problem. *Proceedings of the Section of Sciences of the Koninklijke Nederlandse Akademie van Wetenschappen te Amsterdam* **49(7)**, 758–764.
- [13] Buijs, J. C.; Van Dongen, B. F.; van Der Aalst, W. M. (2012). On the role of fitness, precision, generalization and simplicity in process discovery. In: *OTM Confederated International Conferences "On the Move to Meaningful Internet Systems"*. Springer, pp. 305–322.
- [14] Cadez, I.; Heckerman, D.; Meek, C.; Smyth, P.; White, S. (2000). Visualization of navigation patterns on a web site using model-based clustering. In: *Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining*. pp. 280–284.
- [15] Chierichetti, F.; Kumar, R.; Raghavan, P.; Sarlos, T. (2012). Are web users really markovian? In: *Proceedings of the 21st international conference on World Wide Web*. pp. 609–618.

- [16] Cleary, J.; Witten, I. (1984). Data compression using adaptive coding and partial string matching. *IEEE transactions on Communications* **32(4)**, 396–402.
- [17] Deshpande, M.; Karypis, G. (2004). Selective markov models for predicting web page accesses. *ACM transactions on internet technology (TOIT)* **4(2)**, 163–184.
- [18] Fiedler, M. (1973). Algebraic connectivity of graphs. *Czechoslovak mathematical journal* **23(2)**, 298–305.
- [19] Fournier-Viger, P.; Gomariz, A.; Gueniche, T.; Soltani, A.; Wu, C.-W.; Tseng, V. S. (2014). SPMF: a Java open-source pattern mining library. *The Journal of Machine Learning Research* **15(1)**, 3389–3393.
- [20] Fournier-Viger, P.; Lin, J. C.-W.; Kiran, R. U.; Koh, Y. S.; Thomas, R. (2017). A survey of sequential pattern mining. *Data Science and Pattern Recognition* **1(1)**, 54–77.
- [21] Frias-Martinez, E.; Karamcheti, V. (2002). A prediction model for user access sequences. In: *WEBKDD Workshop: Web Mining for Usage Patterns and User Profiles*.
- [22] Gueniche, T.; Fournier-Viger, P.; Raman, R.; Tseng, V. S. (2015). CPT+: Decreasing the time/space complexity of the Compact Prediction Tree. In: *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. pp. 625–636.
- [23] Gündüz, Ş.; Özsu, M. T. (2003). A web page prediction model based on click-stream tree representation of user behavior. In: *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*. pp. 535–540.
- [24] Kivela, M.; Arenas, A.; Barthelemy, M.; Gleeson, J. P.; Moreno, Y.; Porter, M. A. (2014). Multilayer networks. *Journal of Complex Networks* **2(3)**, 203–271.
- [25] Kraus, M.; Feuerriegel, S. (2019). Personalized Purchase Prediction of Market Baskets with Wasserstein-Based Sequence Matching. In: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. pp. 2643–2652.
- [26] Laird, P.; Saul, R. (1994). Discrete sequence prediction and its applications. *Machine learning* **15(1)**, 43–68.
- [27] Lambiotte, R.; Rosvall, M.; Scholtes, I. (2019). From networks to optimal higher-order models of complex systems. *Nature Physics* **15**, 313–320.
- [28] LaRock, T.; Nanumyan, V.; Scholtes, I.; Casiraghi, G.; Eliassi-Rad, T.; Schweitzer, F. (2019). HYPAs: Efficient Detection of Path Anomalies in Time Series Data on Networks.
- [29] Leemans, S. J.; Tax, N.; ter Hofstede, A. H. (2018). Indulpet miner: Combining discovery algorithms. In: *OTM Confederated International Conferences "On the Move to Meaningful Internet Systems"*. Springer, pp. 97–115.
- [30] for London, T. (2014). Rolling Origin and Destination Survey (RODS) database.

- [31] Mellor, A. (2019). Event Graphs: Advances and Applications of Second-order Time-unfolded Temporal Network Models. *Advances in Complex Systems* **22(03)**, 1950006.
- [32] Padmanabhan, V. N.; Mogul, J. C. (1996). Using predictive prefetching to improve world wide web latency. *ACM SIGCOMM Computer Communication Review* **26(3)**, 22–36.
- [33] Pitkow, J.; Pirolli, P. (1999). Mining longest repeating subsequences to predict world wide web surfing. In: *UsENIX symp. on Internet Technologies and systems*. p. 1.
- [34] Rosvall, M.; Bergstrom, C. T. (2008). Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences* **105(4)**, 1118–1123.
- [35] Rosvall, M.; Esquivel, A. V.; Lancichinetti, A.; West, J. D.; Lambiotte, R. (2014). Memory in network flows and its effects on spreading dynamics and community detection. *Nature communications* **5**, 4630.
- [36] Saebi, M.; Ciampaglia, G. L.; Kaplan, L. M.; Chawla, N. V. (2019). HONEM: Network Embedding Using Higher-Order Patterns in Sequential Data. *arXiv preprint arXiv:1908.05387* .
- [37] Scholtes, I. (2017). 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*. ACM, pp. 1037–1046.
- [38] Scholtes, I.; Wider, N.; Garas, A. (2016). Higher-order aggregate networks in the analysis of temporal networks: path structures and centralities. *The European Physical Journal B* **89(3)**, 61.
- [39] Scholtes, I.; Wider, N.; Pfitzner, R.; Garas, A.; Tessone, C. J.; Schweitzer, F. (2014). Causality-driven slow-down and speed-up of diffusion in non-Markovian temporal networks. *Nature communications* **5**, 5024.
- [40] Singer, P.; Helic, D.; Taraghi, B.; Strohmaier, M. (2014). Detecting memory and structure in human navigation patterns using markov chain models of varying order. *PloS one* **9(7)**, e102070.
- [41] Strelhoff, C. C.; Crutchfield, J. P.; Hübler, A. W. (2007). Inferring Markov chains: Bayesian estimation, model comparison, entropy rate, and out-of-class modeling. *Phys. Rev. E* **76**, 011106.
- [42] Tax, N.; Teinmaa, I.; van Zelst, S. J. (2018). An interdisciplinary comparison of sequence modeling methods for next-element prediction. *arXiv preprint arXiv:1811.00062* .
- [43] Torres, L.; Blevins, A. S.; Bassett, D. S.; Eliassi-Rad, T. (2020). The why, how, and when of representations for complex systems. *arXiv preprint arXiv:2006.02870* .
- [44] TransStat, R. (2014). Origin and Destination Survey database.
- [45] Weijters, A.; Ribeiro, J. (2011). Flexible heuristics miner (FHM). In: *2011 IEEE symposium on computational intelligence and data mining (CIDM)*. IEEE, pp. 310–317.
- [46] West, R.; Leskovec, J. (2012). Human wayfinding in information networks.
- [47] Xu, J.; Wickramaratne, T. L.; Chawla, N. V. (2016). Representing higher-order dependencies in networks. *Science Advances* .