

Interesting Paths in the Mapper

Ananth Kalyanaraman*, Methun Kamruzzaman† and Bala Krishnamoorthy‡

Abstract

Given a high dimensional point cloud of data with functions defined on the points, the Mapper produces a compact summary in the form of a simplicial complex connecting the points. This summary offers insightful data visualizations, which have been employed in applications to identify subsets of points, i.e., subpopulations, with interesting properties. These subpopulations typically appear as long paths, flares (i.e., branching paths), or loops in the Mapper.

We study the problem of quantifying the interestingness of subpopulations in a given Mapper. First, we create a weighted directed graph $G = (V, E)$ using the 1-skeleton of the Mapper. We use the average values at the vertices (i.e., clusters) of the target function (i.e., a dependent variable) to direct the edges from low to high values. We set the difference between the average values at the vertices (high–low) as the weight of the edge. Covariation of the remaining h functions (i.e., independent variables) is captured by a h -bit binary signature assigned to the edge. An *interesting path* in G is a directed path whose edges all have the same signature. Further, we define the interestingness score of such a path as a sum of its edge weights multiplied by a nonlinear function of their corresponding ranks, i.e., the depths of the edges along the path. The goal is to value more the contribution from an edge deep in the path than that from a similar edge which appears at the start.

Second, we study three optimization problems on this graph G to quantify interesting subpopulations. In the problem MAX-IP, the goal is to find the most interesting path in G , i.e., an interesting path with the maximum interestingness score. We show that MAX-IP is NP-complete. For the case where G is a directed acyclic graph (DAG), which could be a typical setting in many applications, we show that MAX-IP can be solved in polynomial time—in $O(mnd_{in})$ time and $O(mn)$ space, where m , n , and d_{in} are the numbers of edges, vertices, and the maximum indegree of a vertex in G , respectively.

In the more general problem IP, the goal is to find a collection of interesting paths that are edge-disjoint, and the sum of interestingness scores of all paths is maximized. We also study a variant of IP termed k -IP, where the goal is to identify a collection of edge-disjoint interesting paths each with k edges, and the total interestingness score of all paths is maximized. While k -IP can be solved in polynomial time for $k \leq 2$, we show k -IP is NP-complete for $k \geq 3$ even when G is a DAG. We develop heuristics for IP and k -IP on DAGs, which use the algorithm for MAX-IP on DAGs as a subroutine, and run in $O(mnd_{in})$ and $O(mkd_{in})$ time for IP and k -IP, respectively.

1 Introduction

Data sets from many applications come in the form of point clouds often in high dimensions along with multiple functions defined on these points. Topological data analysis (TDA) has emerged in the past two decades as a new field whose goal is to summarize such complex data sets, and facilitate the understanding of their underlying topological and geometric structure. In this paper, we focus on *Mapper* [22], a TDA method that has gained significant traction across various application domains [1, 8, 10, 12, 13, 18, 19, 20, 21, 24].

*School of Electrical Engineering and Computer Science, Washington State University, Pullman, WA, 99164, USA; ananth@wsu.edu

†School of Electrical Engineering and Computer Science, Washington State University, Pullman, WA, 99164, USA; md.kamruzzaman@wsu.edu

‡Department of Mathematics and Statistics, Washington State University, Vancouver, WA, 98686, USA; kbala@wsu.edu

Starting from a point cloud X (typically sampled from a metric space), the Mapper studies the topology of the sublevel sets of a *filter function* $f : X \rightarrow Z \subset \mathbb{R}$. Starting with a cover \mathcal{Z} of Z , the Mapper obtains a cover of the domain X by pulling back \mathcal{Z} through f . This pullback cover is then refined into a connected cover by splitting each of its elements into various clusters using a clustering algorithm that employs another function g defined on X . A compact representation of the data set, also termed Mapper, is obtained by taking the nerve of this connected cover—this is a simplicial complex with one vertex per each cluster, one edge per pair of intersecting clusters, and one p -simplex per non-empty $(p+1)$ -fold intersection in general. The method can naturally consider multiple filter functions $f_i : X \rightarrow Z_i$, with covers \mathcal{Z}_i jointly pulled back to obtain the cover of X . Equivalently, one could consider them together as a *single* vector-valued filter function $\mathbf{f} : X \rightarrow Z \subset \mathbb{R}^h$ for $h \geq 2$.

The Mapper has been used in a growing number of applications from diverse domains recently, ranging from medicine [8, 12, 18, 19, 20, 21, 24] to agricultural biotechnology [10] to basketball player profiles [1] to voting patterns [13]. It is also the main engine in the data analytics software platform of the firm Ayasdi. The key to all these success stories is the ability of Mapper to identify subsets of X , i.e., *subpopulations*, that behave distinctly from the rest of the points. In fact, this feature of Mapper distinguishes it from many traditional data analysis techniques based on, e.g., machine learning, where the goal is usually to identify patterns valid for the entire data set.

Several researchers have recently studied mathematical and foundational aspects of Mapper (see Section 1.2). These results have enabled robust ways to build *one* representative Mapper for any given data set. One could follow the work of Carrière et al. [2], or identify parameters corresponding to a stable range in the persistence diagram of the multiscale mapper [4]. While a collection of mappers built at multiple scales, e.g., the multiscale mapper [4], could provide a more detailed representation of the data, efficient summarization of the entire representation as well as the extraction of insights relevant to the application still remain challenging. Hence working with a single Mapper could be considered the desirable setting from the point of view of most applications.

At the same time, many applications demand more precise quantification of the interesting features in the selected Mapper, as well as to track the corresponding subpopulations as they evolve along such features. In plant phenomics [10], for instance, we are interested in identifying specific varieties (i.e., genotypes) of a crop that show resilient growth rates as several environmental factors vary in specific ways during the entire growing season. Each such subpopulation with the associated variation would suggest a testable hypothesis for the practitioner. For this purpose, it is also desirable to rank these subpopulations in terms of their “interestingness” to the practitioner.

1.1 Our contributions

We propose a framework for quantifying the interestingness of subpopulations in a given Mapper. For the input point cloud X , we assume the Mapper is constructed with h filter functions $f_i : X \rightarrow Z_i$ that represent independent variables, and a target function (i.e., a dependent variable) $g : X \rightarrow Z$. The Mapper could be a high-dimensional simplicial complex depending on the choice of covers \mathcal{Z}_i for Z_i .

Formulation: We create a weighted directed graph $G = (V, E)$ using the 1-skeleton of the Mapper. We use the average values of g at the vertices (i.e., clusters) to direct the edges from low to high values. We set the difference between the average values at the vertices (high–low) as the weight of the edge. Covariation of the h functions f_i is captured by a h -bit binary signature assigned to the edge. We define an *interesting path* in G as a directed path whose edges all have the same signature (all references to a “path” in this paper imply a simple path, i.e., no vertices are repeated). Further, we define the interestingness score of such a path as a sum of its edge weights multiplied by a nonlinear function of their corresponding ranks, i.e., the depths of the edges along the path. The goal is to value more the contribution from an edge deep in the path than that from a similar edge which appears at the start.

Theoretical results: We study three optimization problems on this graph G to quantify interesting subpopulations. In the problem MAX-IP, the goal is to find the most interesting path in G , i.e., an interesting path with the maximum interestingness score. We show that MAX-IP is NP-complete. For the special case where G is a directed acyclic graph (DAG), which could be a typical setting in applications, we show that MAX-IP can be solved in polynomial time—in $O(m\delta_{\max}d_{\text{in}})$ time and $O(mn)$ space where m and n are the number of edges and vertices respectively, and δ_{\max} and d_{in} are the diameter of G and the maximum indegree of any vertex in G respectively. Note that $d_{\text{in}} < n$ and $\delta_{\max} < n$ (as G is a DAG).

In the more general problem IP, the goal is to find a collection of interesting paths such that these paths form an exact cover of E and the overall sum of interestingness scores of all paths is maximum. The collection of paths identified by IP could include some short ones in terms of number of edges. Hence we study also a variant of IP termed k -IP, where the goal is to identify a collection of interesting paths each with k edges for a given number k , an edge in E is part of at most one such path, and the total interestingness score of all paths is maximum. While k -IP can be solved in polynomial time for $k \leq 2$, we show k -IP is NP-complete for $k \geq 3$. Further, we show that k -IP remains NP-complete for $k \geq 3$ even for the case when G is a DAG. Finally, we develop heuristics for IP and k -IP on DAGs, which use the algorithm for MAX-IP on DAGs as a subroutine, and run in $O(mnd_{\text{in}})$ and $O(mkd_{\text{in}})$ time for IP and k -IP, respectively.

Software and application: In this paper, we focus primarily on the theoretical aspects of interesting paths. However, we are also actively developing and maintaining an open source software repository that integrates all our ongoing implementations, and their experimental evaluations and applications. Section 6 briefly reports on this ongoing effort, including our use of plant phenomics [9] as a novel application domain for test, validation, and discovery.

1.2 Related work

In most previous applications of Mapper [1, 8, 12, 13, 18, 19, 20, 21, 24], interesting subpopulations are characterized by features (paths, flares, loops) identified in a visual manner. As far as we are aware, our work proposes the first approach to rigorously quantify the interesting features, and to rank them in terms of their interestingness. The works of Carrière et al. [2, 3] present a rigorous theoretical framework for 1-dimensional Mapper, where the features are identified as points in an extended persistence diagram. But this line of work does not address the relative importance of the features in the context of the application generating the data. Our work can be considered as a post processing of the Mapper identified by the methods of Carrière et al. While our framework can naturally consider multiple filter functions for a given Mapper, we do not address the stability of the interesting paths identified.

The interesting paths problems we study are related to the class of nonlinear shortest path and minimum cost flow problems previously investigated. Non-additive shortest paths have been studied [25], and the more general minimum concave cost network flow problem has been shown to be NP-complete [7, 27]. At the same time, versions of shortest path or minimum cost flow problems where the contribution of an edge depends nonlinearly on its position or depth in the path appear to have not received much attention. Hence the specific problems we study should be of independent interest as a new class of nonlinear longest path (equivalently, shortest path) problems.

2 Methods

We refer the reader to the original paper by Singh et al. [22] for background on Mapper, and recent other work [2, 3, 4] for related constructions. For our purposes, we start with a single Mapper M that is a possibly high dimensional simplicial complex constructed from a point cloud X using filter functions $f_i : X \rightarrow Z_i \subset \mathbb{R}$ for $i = 1, \dots, h$ and another function $g : X \rightarrow \mathbb{R}$. In the setting of a typical application, g could represent

a dependent variable whose relationship with the independent variables represented by f_i is of interest. We assume g is used for clustering within the mapper framework.

2.1 Interesting paths and interestingness scores

Each vertex in M represents a cluster of points from X that have similar values of function g , the dependent variable. An edge in M connects two such clusters containing a non-empty intersection of points. By definition, each edge in M connects clusters belonging to distinct elements of the pullback cover of X , and hence the corresponding values of the filter functions f_i also change when moving along the edge. Therefore, by following a trail of vertices (i.e., clusters) whose average g values are monotonically varying, we can capture subpopulations that gradually or abruptly alter their behavior as measured by g under continuously changing filter intervals. For instance, a plant scientist interested in crop resilience may seek to identify a subset of crop individuals/varieties that exhibit sustained or accelerated growth rates (g) despite potentially adverse fluctuations in temperature (f_1) and humidity (f_2). We formulate the problem of identifying such subpopulations as that of finding interesting edge-disjoint paths in a directed graph.

2.1.1 Graph Formulation

We construct a weighted directed graph $G = (V, E)$ representation of the 1-skeleton of M along with some additional information. Let $n = |V|$ and $m = |E|$ denote the numbers of vertices and edges in G , respectively. We set V as the set of vertices (0-simplices) of M , and E as the set of edges (1-simplices) of M . We assign directions and weights to the edges as follows. Each vertex $u \in V$ denotes a subset of points from X that constitute a partial cluster. We denote this subset as $X(u)$. We let $g(u)$ and $f_i(u)$ denote the average values of the clustering function g (dependent variable) and the filter function f_i , respectively, for all points in u :

$$g(u) = \frac{\sum_{x \in X(u)} g(x)}{|X(u)|} \quad \text{and} \quad f_i(u) = \frac{\sum_{x \in X(u)} f_i(x)}{|X(u)|}, \quad i = 1, \dots, h.$$

We let $\omega(u)$ represent the weight of a vertex u . In this paper, we set $\omega(u)$ to be equal to $g(u)$. For an edge $e = (u, v)$ in E , we assign as its weight as $\omega(e) = |\omega(u) - \omega(v)| = |g(u) - g(v)|$. Notice $\omega(e) \geq 0$ for all edges e in G .

Each edge is also associated with a direction, determined by one of the two rules illustrated in Figure 1. The simpler rule, termed **Rule a**, directs the edge from the vertex with lower weight to the vertex with higher weight (“a” for *ascending* weight). If the vertex weights are equal, one of the two directions is chosen arbitrarily. The other rule, termed **Rule b**, handles differently the case where the weights of the two vertices are close, as defined by a cutoff $\tau \geq 0$. In this case, the edge is allowed to be bidirectional (“b” for *bidirectional*), i.e., both the forward and backward edges are added with the same weight. If the weights of the vertices are not close, **Rule a** is followed. The latter relaxed scheme is motivated toward capturing more robust interesting paths in practice.

Note that if only **Rule a** is followed, the resulting graph is guaranteed to be a directed acyclic graph (DAG). On the other hand, if the relaxed scheme **Rule b** is followed, then the resulting graph is directed, i.e., it could potentially be cyclic. Consequently, the optimization problems we address in this paper (see Section 2.1.2) will be presented for both DAG and directed graph inputs.

We assign a h -bit binary signature $\text{Sig}(e) = b_1 b_2 \dots b_h$ to the oriented edge $e = (u, v)$ (i.e., $e : u \rightarrow v$) to capture the covariation of g and the filter functions f_i . We set $b_i = 1$ if $f_i(u) \leq f_i(v)$, and $b_i = 0$ otherwise. If the edge is bidirected, then the signature is used as a “wildcard”—the signature is not predetermined, and is chosen to match any candidate signature as determined by the interesting path detection algorithm.



Rule a: Direct edge from u to v if $\omega(v) > \omega(u)$, and from v to u if $\omega(v) < \omega(u)$. If the weights are equal, then one of the two directions is chosen arbitrarily.

Rule b: Two opposite directed edges between nodes u and v if $|\omega(v) - \omega(u)| \leq \tau$, and directed as per Rule a otherwise.

Figure 1: Two rules for directing an edge, based on vertex weights and a user-defined cutoff (τ).

Definition 2.1. An *interesting k -path* for a given k with $1 \leq k \leq n - 1$ is a directed path $P = [e_{i_1}, \dots, e_{i_k}]$ of k edges in G , such that $\text{Sig}(e_r)$ is identical for all $r = i_1, \dots, i_k$. An *interesting path* is a path of arbitrary length in the interval $[1, n - 1]$.

Note that in the above definition, a directed path is one where all edges in the path have the same direction. Hence we have the flexibility to use a bidirected edge as part of a directed path either in the forward or reverse direction, but not both.

Definition 2.2. Given an interesting k -path $P = [e_{i_1}, \dots, e_{i_k}]$ in G as specified in Definition 2.1, we define its *interestingness score* as follows.

$$\mathcal{I}(P) = \sum_{r=1}^k \omega(e_{i_r}) \times \log(1 + r). \quad (1)$$

In particular, the contribution of an edge $e \in P$ to $\mathcal{I}(P)$ is set to $\omega(e) \times \log(1 + \text{rank}(e, P))$, where $\text{rank}(e, P)$ is the rank or order of edge e as it appears in P .

Intuitively, we use the rank of an edge as an inflation factor for its weight—the later an edge appears in the path, the more its weight will count toward the interestingness of the path. This logic incentivizes the growth of long paths. The log function, on the other hand, helps temper this growth in terms of number of edges. Inflation of weights for edges that appear later in the path is motivated by the potential interpretation of interesting paths in the context of real world applications. For instance, while analyzing plant phenomics data sets [10, 14], we expect plants to show accelerated growth spurts later in the growth season. Plants showing such spurts later in the season are potentially more interesting to the practitioner than ones showing a steady growth rate throughout the season.

Remark 2.3. The above framework can be modified easily to characterize robust interesting paths, where the signature matching condition is relaxed—for instance, $b_i = 1$ if $0.9f_i(v) \leq f_i(u) \leq f_i(v)$.

Remark 2.4. While we assume g and f_i are functions from X to \mathbb{R} , our framework could handle more general functions as well. If some f_i is a vector-valued function, for instance, we could first compute pairwise distances of the points in X using f_i , and then assign to each point in X its average distance to all other points as a “surrogate” function.

Remark 2.5. The framework applies without change to cases where g is used along with other functions to cluster. In fact, g could be used also as a filter function as long as it is used for clustering.

2.1.2 Optimization Problems

We now present multiple optimization problems with the broader goal of identifying interesting path(s) that maximize interestingness score(s).

MAX-IP: Find an interesting path P in G such that $\mathcal{I}(P)$ is maximized.

k -IP: For a given k between 1 and $n - 1$, find a collection \mathcal{P} of interesting k -paths such that each $e \in E$ is part of at most one $P \in \mathcal{P}$, and the total interestingness score $\mathcal{I}(\mathcal{P}) = \sum_{P \in \mathcal{P}} \mathcal{I}(P)$ is maximized.

IP: Find a collection \mathcal{P} of interesting paths in G such that the total interestingness score $\mathcal{I}(\mathcal{P}) = \sum_{P \in \mathcal{P}} \mathcal{I}(P)$ is maximized (\mathcal{P} will exactly cover E , i.e., each $e \in E$ is part of exactly one $P \in \mathcal{P}$).

Both IP and k -IP produce edge-disjoint collections of interesting paths. In IP, every edge in G is part of an interesting path in \mathcal{P} . But this setting might include several short (in number of edges) interesting paths. In k -IP, each interesting path found has exactly k edges, and some edges in E might not be part of any interesting k -path in \mathcal{P} . Hence the paths identified by k -IP are likely to be more meaningful in practice.

Remark 2.6. The $\log(1 + \text{rank})$ factor in the interestingness score in Equation (1) makes each of the above optimization problems nonlinear. At the same time, the type of nonlinearity introduced here is distinct from the ones studied in the literature, e.g., in non-additive shortest paths [25], or in minimum concave cost flow [7, 27]. Hence these problems form a new class of nonlinear longest (equivalently, shortest) path problems, which would be of interest independent of their application in the context of the Mapper and TDA.

Remark 2.7. Our directed graph formulation (in Section 2.1.1) with signatures determined by h filter functions could also be considered equivalently as the computation of the coboundary of the filter functions seen as a 0-cochain with coefficients in \mathbb{R}^h [17]. Further, under appropriate assumptions on the filter functions being smooth, candidates for interesting paths could be seen as flows in a gradient field [26, §14]. Under these assumptions, one could argue that the graph constructed will necessarily be a DAG, and results from Morse theory [15, 16] would also apply. But we are not assuming the functions involved are necessarily smooth. More importantly, the relaxed Rule b (illustrated in Figure 1) creates bidirectional edges, which are more appropriate in the context of real world applications. Further, the $\log(1 + \text{rank})$ factor used in defining our interestingness score (in Equation 1) is not captured by default approaches for maximal flows in gradient fields or by Morse theory.

3 The MAX-IP Problem

The goal of MAX-IP is to identify an interesting path with the maximum interestingness score. We show MAX-IP is NP-complete on directed graphs, but is in P on directed acyclic graphs (DAGs).

3.1 MAX-IP on directed graphs

In the decision version of MAX-IP termed MAX-IPD, we are given a directed graph $G = (V, E)$ with edge weights $\omega(e) \geq 0$ and signatures $\text{Sig}(e)$ for $e \in E$ and a target score $s_0 \geq 0$. The goal is to determine if there exists an interesting path P in G whose interestingness score $\mathcal{I}(P) \geq s_0$.

Lemma 3.1. MAX-IPD on directed graph $G = (V, E)$ is NP-complete.

Proof. Given a path P in G , we can verify that it is a directed simple path with all the edges of the same signature, compute its interestingness score $\mathcal{I}(P)$ using Equation (1), and compare it with s_0 —all in polynomial time. Hence MAX-IPD is in NP.

We reduce the problem of checking if a directed graph has a directed Hamiltonian cycle (DIRHC) to MAX-IPD. DIRHC is one of the 21 NP-complete problems originally introduced by Karp [11]. Given an instance $G = (V, E)$ of DIRHC with $|V| = n$, we construct an instance of MAX-IPD on a directed graph $G' = (V', E')$ as follows. We replace an arbitrary vertex $v \in V$ by two vertices v' and v'' , i.e.,

$V' = (V \setminus \{v\}) \cup \{v', v''\}$ and $|V'| = n + 1$. Each $(v, w) \in E$ is replaced by (v', w) in E' and each $(u, v) \in E$ is replaced by (u, v'') in E' . All other edges in E are included in E' without changes. All edges in E' are assigned unit weights and identical signatures, and we set $s_0 = \log((n + 1)!)$.

We claim that G has a directed Hamiltonian cycle C if and only if there exists an interesting path P' in G' with interestingness score $\mathcal{I}(P') = s_0$. Let G have a directed Hamiltonian cycle C . Then C must have n edges by definition, and C visits (i.e., enters and leaves) each vertex in V exactly once. Hence there must exist edges (v, w) and (u, v) in C . We construct the interesting path P' in G' using $(v', w), (u, v'')$, and the remaining $(n - 2)$ edges in C . Thus P' is a directed path in G' with n edges. Further, since all edges in E' have unit weights and identical signatures, it is clear from Equation (1) that P' is indeed an interesting path in G' with $\mathcal{I}(P') = s_0$.

Conversely, let P' be an interesting path in G' with $\mathcal{I}(P') = s_0$ (notice that $\mathcal{I}(P') > s_0$ is not possible, as nodes are not allowed to be repeated in P'). Since P is an interesting path, it visits (i.e., enters and/or leaves) any vertex in V' at most once. Since all edges in E' have unit weights and identical signatures, and by the definition of interestingness score in Equation (1), it is clear that P' must have n edges. Hence P' must start with an edge (v', w) and end with an edge (u, v'') . Then the directed cycle C in G defined by the edges $(v, w), (u, v)$, and the remaining $(n - 2)$ edges in P' is Hamiltonian. Hence MAX-IPD on directed graphs is NP-complete. \square

3.2 MAX-IP on directed acyclic graphs

Lemma 3.2. MAX-IP on a directed acyclic graph $G = (V, E)$ is in P .

Proof. We present a polynomial time algorithm for MAX-IP on a DAG (as proof of Lemma 3.2). The input is a DAG $G = (V, E)$ with n vertices and m edges, with edge weights $\omega(e) \geq 0$ and signatures $\text{Sig}(e)$ for all $e \in E$. The output is an interesting path P^* which has the maximum interestingness score in G . We use dynamic programming, with the forward phase computing $\mathcal{I}(P^*)$ and the backtracking procedure reconstructing a corresponding P^* .

Let $T(i, j)$ denote the score of a maximum interesting path of length j edges ending at edge e_i for $i \in [1, m]$. Since an interesting path could be of length at most $(n - 1)$, we have $j \in [1, n - 1]$. Therefore the values in the recurrence can be maintained in a 2-dimensional table of size $m \times (n - 1)$, as illustrated in Figure 2. The algorithm has three steps:

- **Initialization:** $T(i, 1) = \omega(e_i) \times \log(2)$, where $1 \leq i \leq m$.
- **Recurrence:** For an edge $e = (u, v) \in E$, we define a *predecessor edge* of e as any edge $e' \in E$ of the form $e' = (w, u)$ and $\text{Sig}(e') = \text{Sig}(e)$. Let $\text{Pred}(e)$ denote the set of all predecessor edges of e . Note that $\text{Pred}(e)$ can be possibly empty. We define the recurrence for $T(i, j)$ as follows.

$$T(i, j) = \max_{e_{i'} \in \text{Pred}(e_i)} \{T(i', j - 1) + \omega(e_i) \times \log(1 + j)\} \quad (2)$$

- **Output:** We report the score that is maximum in the entire table. A corresponding optimal path P^* can be obtained by backtracking from that cell to the first column.

Proof of correctness: Any interesting path in G can be at most $n - 1$ edges long. As a particular edge could appear anywhere along such a path, its rank can range between 1 and $n - 1$. Hence the $m \times (n - 1)$ recurrence table T sufficiently captures all possibilities for each edge in E . The following key observation completes the proof. Let $P^*(i, j)$ denote an optimal scoring path, if one exists, of length $j \in [1, n - 1]$ ending at edge $e_i \in E$. If $P^*(i, j)$ exists and if $j > 1$, then there should also exist $P^*(i', j - 1)$ where $i' \in \text{Pred}(e_i)$. Furthermore, the edge e_i could not have appeared in $P^*(i', j - 1)$ because G is acyclic.

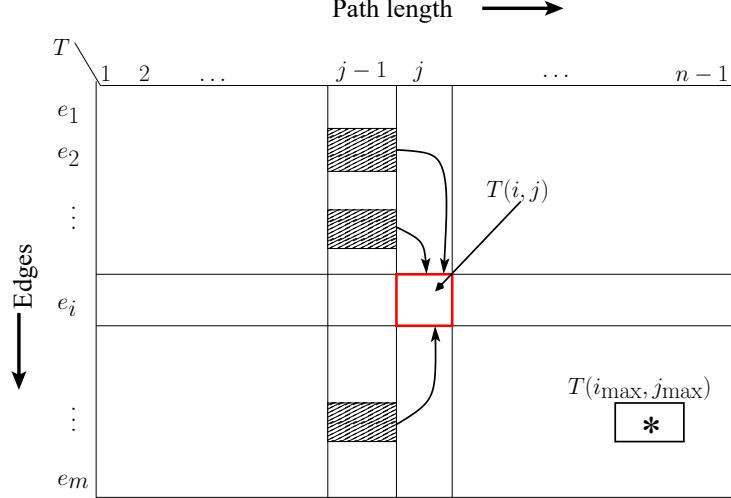


Figure 2: Table $T(i, j)$ for the MAX-IP algorithm.

Therefore, due to the edge-disjoint nature of $P^*(i', j - 1)$ and the remainder of $P^*(i, j)$ (which is e_i), the principle of optimality is preserved—i.e., the maximum operator in Equation (2) is guaranteed to ensure optimality of $T(i, j)$.

Complexity analysis: The above dynamic programming algorithm can be implemented to run in $O(mn)$ space and a worst-case time complexity of $O(mnd_{\text{in}})$, where d_{in} denotes the maximum indegree of any vertex in V . \square

3.2.1 Algorithmic improvements

The above dynamic programming algorithm for MAX-IP for DAGs can be implemented to run in space and time smaller in practice than the worst case limits suggested above. First, we note that computing the full table T is likely to be wasteful, as it is likely to be sparse in practice. The sparsity of T follows from the observation that an interesting path of length j ending at edge e_i can exist only if there exists at least one other interesting path of length $j - 1$ ending at one of e_i 's predecessor edges. We can exploit this property by designing an iterative implementation as follows.

Instead of storing the entire table T , we store only the rows (edges), and introduce columns on a “need basis” by maintaining a dynamic list $L(e_i)$ of column indices for each edge e_i .

- S1) Initially, we assign $L(e_i) = \{1\}$, as each edge is guaranteed to be in an interesting path of length at least 1 (the path consisting of the edge by itself).
- S2) In general, the algorithm performs multiple iterations; within each iteration, we visit and update the dynamic lists for all edges in E as follows. For every edge $e_{i'} \in \text{Pred}(e_i)$, $L(e_i) = L(e_i) \cup \{\ell + 1 \mid \ell \in L(e_{i'})\}$. The algorithm iterates until there is no further change in the lists for any of the edges.

The number of iterations in the above implementation can be bounded by the length of the longest path in the DAG (i.e., the diameter δ_{max}), which is less than n . Also, we implement the list update from predecessors to successors such that each edge is visited only a constant number of times (despite the varying products of in- and out-degrees at different vertices). To this end, we implement the update in S2 as a two-step process: first, performing a union of all lists from the predecessor edges of the form $(*, v)$ so that the merged lists can be used to update the lists of all the successor edges of the form $(v, *)$. Thus the work in each iteration is bounded by $O(m)$.

Taken together, even in the worst-case scenario of $(\delta_{\max} + 1)$ iterations, the overall time to construct these dynamic lists is $O(m\delta_{\max})$. Furthermore, during the list construction process, if one were to carefully store the predecessor locations using pointers, then the computation of the $T(i, j)$ recurrence in each cell can be executed in time proportional to the number of *non-empty* predecessor values in the table. Overall, this revised algorithm can be implemented to run in time $O(m\delta_{\max}d_{\text{in}})$, and in space proportional to the number of non-zero values in the matrix.

Further, the above implementation is also inherently parallel since the list value at an edge in the current iteration depends only on the list values of its predecessors from the previous iteration.

4 The k -IP Problem

The goal of k -IP for $k \leq n - 1$ is to find a set \mathcal{P} of edge-disjoint interesting k -paths such that the sum of their interestingness scores is maximized. We show that k -IP on directed graphs can be solved in polynomial time for $k \leq 2$. On the other hand, we show that k -IP on a directed acyclic graph is NP-Complete for $k \geq 3$ (see Theorem 4.2).

The smallest value of k for which k -IP is nontrivial is 2. We can solve 2-IP as a weighted matching problem.

Lemma 4.1. *k -IPD on directed graph $G = (V, E)$ is in P for $k \leq 2$.*

Proof. The case of $k = 1$ turns out to be trivial. An optimal solution for 1-IP is obtained by taking \mathcal{P} as a collection of m interesting 1-paths each comprised of a single edge. These 1-paths are edge-disjoint by definition, and the need to compare signatures within a path does not arise. Since all edge weights $\omega \geq 0$, the total interestingness score $\mathcal{I}(\mathcal{P})$ is guaranteed to be maximum. This optimal solution is unique when $\omega(e) > 0$ for all edges $e \in E$.

We model the 2-IP problem ($k = 2$) as an equivalent weighted matching problem on an undirected graph $G' = (V', E')$, which we construct as follows. We include a vertex $i \in V'$ for each edge $e_i \in E$ in the input graph. Hence $|V'| = m$. Whenever edges $\{e_i, e_j\} \in E$ form an interesting 2-path P_{ij} in G , we add the undirected edge $(i, j) \in E'$ with its weight $\omega_{ij} = \mathcal{I}(P_{ij})$ computed using Equation (1). If both interesting paths $P_{ij} = [e_i, e_j]$ and $P_{ji} = [e_j, e_i]$ are possibly formed by a pair of edges $\{e_i, e_j\} \in E$, we set $\omega_{ij} = \max\{\mathcal{I}(P_{ij}), \mathcal{I}(P_{ji})\}$. Notice that $\omega_{ij} \geq 0$ for all edges $(i, j) \in E'$, and $|E'| \leq m(m - 1)/2$. A matching $M' \subseteq E'$ in G' corresponds to a set \mathcal{P} of edge-disjoint interesting 2-paths in G —a vertex $i \in V'$ will be matched with *at most* one other vertex $j \in V'$, and such a match of vertices in V' corresponds to the interesting path $P_{ij} \in \mathcal{P}$ (or $P_{ji} \in \mathcal{P}$, but not both). It follows that a maximum matching in G' corresponds to an optimal solution to 2-IP on the input graph G .

The maximum weighted matching problem on an undirected graph with n vertices and m edges can be solved in strongly polynomial time—e.g., Gabow’s implementation [6] of Edmonds’ algorithm [5] runs in $O(nm + n^2 \log n)$ time. As such, we can solve 2-IP by solving the weighted matching problem on the associated graph G' in $O(m^3)$ time. Hence k -IP is in P for $k \leq 2$. \square

We now consider k -IP for $k \geq 3$ on directed acyclic graphs. To characterize its complexity, we study the decision version of k -IP termed k -IPD, in which we are given a directed acyclic graph $G = (V, E)$ with edge weights $\omega(e) \geq 0$ and signatures $\text{Sig}(e)$ for $e \in E$ and a target score $s_0 \geq 0$. The goal is to determine if there exists a collection \mathcal{P} of edge-disjoint interesting k -paths in G whose total interestingness score $\mathcal{I}(\mathcal{P}) = s_0$.

Theorem 4.2. *k -IPD on a directed acyclic graph $G = (V, E)$ is NP-complete for $k \geq 3$.*

Proof. Given a collection \mathcal{P} of interesting k -paths in a directed acyclic graph $G = (V, E)$, we can verify they are edge-disjoint, each path has k edges, and signatures are identical for all edges in each path, all in

polynomial time. We can compute the interestingness score of each k -path $P \in \mathcal{P}$ using Equation (1) also in polynomial time, and add the $\mathcal{I}(P)$ values to compare with s_0 to check for equality. Hence k -IPD is in NP.

We now reduce the *exact 3-cover* problem (X3C) to 3-IPD. We then show a similar reduction for $k \geq 4$ as well, proving k -IPD is NP-complete for $k \geq 3$. The latter case for general k subsumes the case for $k = 3$. We still present the details for $k = 3$ separately, as this case reveals the structure of the general reduction in an arguably simpler setting. X3C is a version of one of the 21 NP-complete problems originally introduced by Karp [11], and is defined as follows. Given a set X with $|X| = 3q$ elements and a collection \mathcal{C} of 3-element subsets of X with $|\mathcal{C}| = p$, determine if there exists a subset $\mathcal{C}' \subseteq \mathcal{C}$ such that each element of X belongs to exactly one member of \mathcal{C}' . Notice that such an exact cover \mathcal{C}' must necessarily have exactly q members. Also, we assume $p \geq q \geq 3$ (else the instance will be trivial).

Given an instance of X3C, we create a directed acyclic graph $G = (V, E)$ for an instance of 3-IPD as follows. Each element $x \in X$ corresponds to a unique directed edge in G . Corresponding to each 3-element set $\{x, y, z\} \in \mathcal{C}$, we add to G a directed acyclic graph object as shown in Figure 3. The edges corresponding to all $x \in X$ are assigned the large weight $\omega = p$ making them “heavy” edges, while the rest of the edges are all assigned unit weights. Further, we assume $\text{Sig}(e)$ is identical for all edges $e \in E$. The three “V”-shaped 3-paths in the top of Figure 3 are referred to as the x -, y -, and z -paths. Notice that by this construction, G can have at most $|V| \leq 13p$ vertices and $|E| \leq 12p$ edges.

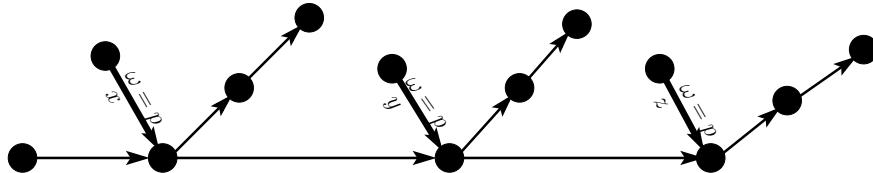


Figure 3: Graph object corresponding to set $\{x, y, z\} \in \mathcal{C}$ in X3C. Thin edges all have $\omega = 1$.

Let $W_{\text{in}} = 3(p \log 2 + \log 3 + \log 4) + (\log 2 + \log 3 + \log 4) = 4 \log 24 + 3(p - 1) \log 2$, and $W_{\text{out}} = 3 \log 24$. From a graph object as shown above, we observe that 4 edge-disjoint interesting 3-paths can be chosen by 3-IPD each with interestingness score W_{in} if and only if the x -, y -, and z -paths are chosen along with *one* other interesting 3-path as shown in Figure 4a. Further, each edge corresponding to an element in X may belong to only one 3-path. Thus, at most q such graph objects may contribute the score of W_{in} to the total interestingness score. The remaining $p - q$ graph objects may contribute a score of at most W_{out} corresponding to the selection of the 3 edge-disjoint interesting 3-paths shown in Figure 4b, which avoid the edges corresponding to any $x \in X$. If q such graph objects do contribute W_{in} each to the total interestingness score, it is clear that the corresponding q triplet elements in \mathcal{C} form an exact 3-cover of X . Further, 3-IPD on G will identify exactly $4q + 3(p - q) = 3p + q$ edge-disjoint interesting 3-paths with a total interestingness score of exactly $qW_{\text{in}} + (p - q)W_{\text{out}} = pW_{\text{out}} + 3(p - 1)q \log 2 + q \log 24$.

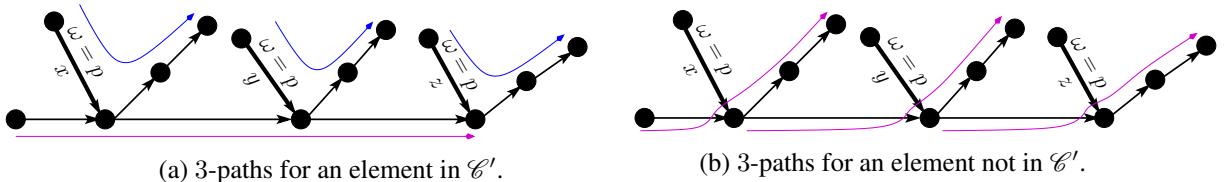


Figure 4: Choices of interesting 3-paths in graph objects for elements in \mathcal{C}' and $\mathcal{C} \setminus \mathcal{C}'$.

Conversely, if X3C has an exact 3-cover \mathcal{C}' , we choose the 4 edge-disjoint interesting 3-paths (recall we assume identical signatures for all edges in G) with interestingness score W_{in} as described above in the

graph object in G corresponding to each of the q 3-element sets in \mathcal{C}' . For the $(p - q)$ 3-element sets in $\mathcal{C} \setminus \mathcal{C}'$, we choose the 3 interesting 3-paths with interestingness score W_{out} each in the corresponding graph objects in G . This collection of p edge-disjoint interesting 3-paths in G will have a total interestingness score of exactly $pW_{\text{out}} + 3(p - 1)q \log 2 + q \log 24$.

Thus X has an exact 3-cover if and only if 3-IPD on G has a target total interestingness score of $s_0 = pW_{\text{out}} + 3(p - 1)q \log 2 + q \log 24$, proving 3-IPD is NP-complete.

We now extend this result to k -IPD for $k \geq 4$. To this end, we reduce the *exact k -cover* problem (XkC) to k -IPD for general $k \geq 4$. The XkC problem is a generalization of X3C, and is defined as follows. Given a set X with $|X| = kq$ elements and a collection \mathcal{C} of k -element subsets of X with $|\mathcal{C}| = p$, determine if there exists a subset $\mathcal{C}' \subseteq \mathcal{C}$ such that each element of X belongs to exactly one member of \mathcal{C}' . Notice that such an exact cover \mathcal{C}' must necessarily have exactly q members. Also, we assume $p \geq q \geq k$ (else the instance will be trivial).

Given an instance of XkC, we create a directed acyclic graph G for an instance of k -IPD as follows. Each element $x \in X$ corresponds to a unique directed edge in G . For each k -element set $\{x_1, x_2, \dots, x_k\} \in \mathcal{C}$, we add to G a corresponding directed acyclic graph object as shown in Figure 5. The edges corresponding to all $x \in X$ are assigned the large weight $\omega = p$ (giving “heavy” edges), while the rest of the edges are all assigned unit weights. Further, we assume $\text{Sig}(e)$ is identical for all edges $e \in E$. The k “V”-shaped k -paths in the top of Figure 5 are referred to as the x_1 -, x_2 -, \dots , x_k -paths. Notice that by this construction, G can have at most $|V| \leq (k(k + 1) + 1)p$ vertices and $|E| \leq (k(k + 1))p$ edges.

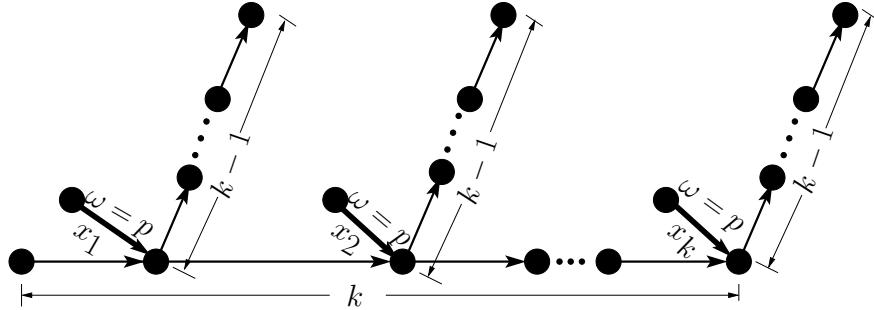


Figure 5: Graph object corresponding to set $\{x_1, x_2, \dots, x_k\} \in \mathcal{C}$ in XkC. Thin edges all have $\omega = 1$.

Let $W_{\text{in}} = k(p \log 2 + \log 3 + \log 4 + \dots + \log(k + 1)) + (\log 2 + \log 3 + \dots + \log(k + 1)) = (k + 1) \log((k + 1)!) + k(p - 1) \log 2$, and $W_{\text{out}} = k \log((k + 1)!)$. From a graph object as shown above, we observe that $k + 1$ edge-disjoint interesting k -paths can be chosen by k -IPD each with interestingness score W_{in} if and only if the x_1 -, x_2 -, \dots , x_k -paths are chosen along with *one* other k -path as shown in Figure 6a. Further, each edge corresponding to an element in X may belong to only one k -path. Thus, at most q such graph objects may contribute the score of W_{in} each to the total interestingness score. The remaining $p - q$ graph objects may contribute a score of at most W_{out} each corresponding to the selection of the k edge-disjoint interesting k -paths shown in Figure 6b, which avoid the edges corresponding to any $x \in X$. If q such graph objects do contribute W_{in} each to the total interestingness score, it is clear that the corresponding q k -tuple elements in \mathcal{C} form an exact k -cover of X . Further, k -IPD on G will identify exactly $(k + 1)q + k(p - q) = kp + q$ edge-disjoint interesting k -paths with a total interestingness score of exactly $qW_{\text{in}} + (p - q)W_{\text{out}} = pW_{\text{out}} + k(p - 1)q \log 2 + q \log((k + 1)!)$.

Conversely, if XkC has an exact k -cover \mathcal{C}' , we choose the $k + 1$ edge-disjoint interesting k -paths (again, we assume identical signatures for all edges in G) with interestingness score W_{in} as described above in the graph object in G corresponding to each of the q k -element sets in \mathcal{C}' . For the $(p - q)$ k -element sets in $\mathcal{C} \setminus \mathcal{C}'$, we choose the k edge-disjoint interesting k -paths with interestingness score W_{out} each in

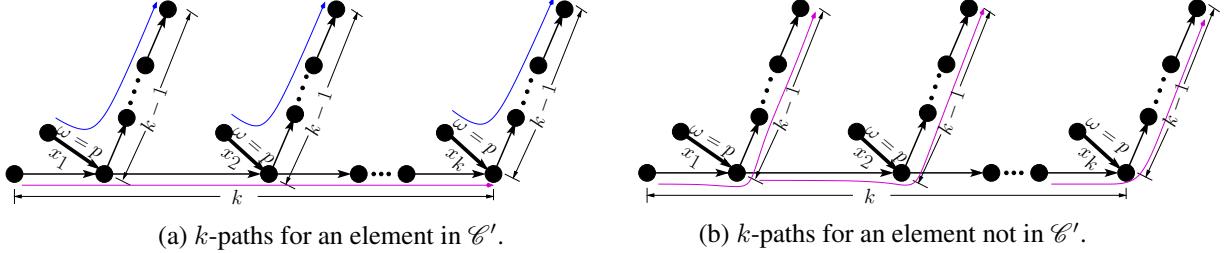


Figure 6: Choices of interesting k -paths in graph objects for elements in \mathcal{C}' and $\mathcal{C} \setminus \mathcal{C}'$.

the corresponding graph objects in G . This collection of p edge-disjoint interesting k -paths in G will have a total interestingness score of exactly $pW_{\text{out}} + k(p-1)q \log 2 + q \log((k+1)!)$.

Thus X has an exact k -cover if and only if k -IPD on G has a target total interestingness score of $s_0 = pW_{\text{out}} + k(p-1)q \log 2 + q \log((k+1)!)$, proving k -IPD is NP-complete for $k \geq 4$. \square

Note that since k -IP is NP-complete for DAGs, it is also NP-complete for directed graphs.

5 The Interesting Paths (IP) Problem

The goal of IP is to find a set \mathcal{P} of edge-disjoint interesting paths of possibly varying lengths (1 to $n-1$) such that the sum of their interestingness scores is maximized. The determination of the precise complexity class for IP is an open problem. However, based on the hardness results for k -IP (Section 4), we conjecture that the IP is also intractable. Here we present an efficient heuristic for IP on DAGs by employing the exact algorithm for MAX-IP on DAGs (in Section 3.2) as a subroutine. We also estimate lower and upper bounds on the maximum total interestingness score of IP (see Section 5.2).

5.1 An efficient heuristic for IP on DAGs

We present a polynomial time heuristic to find a set of edge-disjoint interesting paths \mathcal{P} in a DAG with high total interestingness score. We do not provide any guarantee on the optimality or (approximation) quality of the collection of interesting paths \mathcal{P} .

Our method, termed Algorithm 1, uses a greedy strategy by iteratively calling the exact algorithm for MAX-IP (Section 3.2). The idea is to iteratively detect a maximum interesting path, add it to the working set of solutions, remove all the edges in that path, and recompute MAX-IP on the remaining graph, until there are no more edges left.

Algorithm 1: Greedy Heuristic for IP on DAGs

Input: DAG $G = (V, E)$ with $\omega(e)$, $\text{Sig}(e) \forall e \in E$
Output: A set of edge-disjoint interesting paths \mathcal{P} in G
 $\mathcal{P} = \emptyset$
repeat
 $P \leftarrow$ Compute MAX-IP on $G = (V, E)$ and return a most interesting path
 $\mathcal{P} \leftarrow \mathcal{P} \cup \{P\}$
Remove edges in P from E
until $E = \emptyset$;
return \mathcal{P}

Complexity Analysis: The runtime to compute MAX-IP on $G = (V, E)$ in the first step is $O(m\delta_{\max}d_{\text{in}})$, as described in Section 3.2.1. Recall that δ_{\max} is the diameter of the DAG and d_{in} is the maximum indegree of any vertex in V . Therefore, if we denote p to be the number of iterations (equivalently, the number of interesting paths found), then the overall runtime complexity is $O(pm\delta_{\max}d_{\text{in}})$. However, we expect the performance of the algorithm in practice to be much faster. Note that at least one edge is, and at most m edges are, eliminated in each iteration, thereby implying $1 \leq p \leq m$ here.

Consider the worst case of elimination where one edge is eliminated in each iteration, i.e., $p = m$. The graph must be very sparse in this case, i.e., $m = \Theta(n)$, causing our algorithm for MAX-IP to perform only $O(m)$ work per iteration. Therefore the overall runtime is $O(m^2)$, or equivalently, $O(n^2)$.

On the other hand, consider the case where the number of edges reduces by a constant factor c at every iteration. This setting implies $p = O(\log_c(m))$, while the work performed from one iteration to the next will also continually reduce by a factor of c . Hence the overall runtime can still be bounded by $O(m\delta_{\max}d_{\text{in}})$, the cost of MAX-IP. Further, from an application standpoint, such a greedy iterative approach can be terminated whenever an adequate number of “top” interesting paths are identified.

5.1.1 An efficient heuristic for k -IP on DAGs

The above heuristic for IP can be easily modified to devise a heuristic for k -IP on DAGs. Algorithm 2 summarizes the main steps. The main idea is to modify the exact algorithm for MAX-IP on a DAG such that it initializes a recurrence table of size $m \times k$ (instead of $m \times (n - 1)$), and then use that table to iteratively compute MAX-IP paths. The only constraint here is that each such MAX-IP path should originate from the k -th column during backtracking, so that paths output are guaranteed to be of length k . The runtime is bounded by $O(mrd_{\text{in}})$, where $r = \min\{k, \delta_{\max}\}$.

Algorithm 2: Greedy Heuristic for k -IP on DAGs

Input: DAG $G = (V, E)$ with $\omega(e)$, $\text{Sig}(e) \forall e \in E$
Output: A set of edge-disjoint interesting k -paths \mathcal{P} in G
 Initialize an $m \times k$ table T'
 $\mathcal{P} = \emptyset$
repeat
 $P' \leftarrow$ Compute MAX-IP on $G = (V, E)$ using T' , and return
 a most interesting path *ending in column k*
 $\mathcal{P} \leftarrow \mathcal{P} \cup \{P'\}$
 Remove edges in P' from E
until $E = \emptyset$;
return \mathcal{P}

Remark 5.1. Ideas used in Algorithms 1 and 2 could be combined to develop a heuristic for ATLEAST- k -IP, a modified version of k -IP that seeks to find a collection of interesting paths in G where each path has *at least* k edges. We have implemented this heuristic in our software suite (see Section 6 for details).

5.2 Bounds for IP

Let \mathcal{P}^* represent an optimal set of paths for an instance of IP. We derive upper and lower bounds on its total interestingness score $\mathcal{I}(\mathcal{P}^*)$.

Let $P_{\max}(i)$ denote a maximum interesting path (of arbitrary length) ending at a given edge $e_i \in E$.

Lemma 5.2. $\mathcal{I}(\mathcal{P}^*) \leq \sum_{e_i \in E} \mathcal{I}(P_{\max}(i))$.

Proof. Consider an arbitrary path $P \in \mathcal{P}^*$. We first note that individual paths that are members of an optimal solution (\mathcal{P}^*) for the IP problem can end at any arbitrary *non-source* vertex in $G = (V, E)$ (see Figure 7 for an illustration).

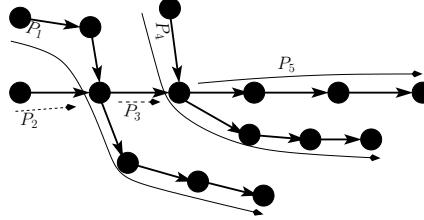


Figure 7: Five interesting paths labeled P_1, \dots, P_5 . The paths P_2 and P_3 end at non-source vertices.

Without loss of generality, let us assume that the input graph contains only vertices with degree at least one (as vertices with degree zero cannot contribute to any interesting path). We consider two sub-cases:

Case A: No two maximum scoring paths ending at two different edges e_i and e_j intersect, i.e., $P_{\max}(i) \cap P_{\max}(j) = \emptyset, \forall i, j \in [1, m]$ and $i \neq j$.

This case can occur only if the number of edges (m) is equal to the number of source vertices. This setting implies that \mathcal{P}^* is comprised of m paths, where each path $P \in \mathcal{P}^*$ is a unique edge $e_i \in E$. Therefore, $\mathcal{I}(\mathcal{P}^*) = \sum_{e_i \in E} \mathcal{I}(P_{\max}(i))$ in this case.

Case B: There exists at least two maximum interesting paths ending at two different edges e_i and e_j that intersect, i.e., $P_{\max}(i) \cap P_{\max}(j) \neq \emptyset$.

This case implies that at least one of these two paths is *not* a member of \mathcal{P}^* (by definition of the IP problem); let this non-member path be $P_{\max}(j)$ without loss of generality. Since all edges are covered by \mathcal{P}^* by definition of IP, there still has to exist an alternative path ending in e_j that is either directly contained in \mathcal{P}^* or contained as a *subpath* of a longer path in \mathcal{P}^* ; let us refer to this alternative path as $P'(j)$. Since $P_{\max}(j)$ is an optimal interesting path ending at edge e_j , $\mathcal{I}(P'(j)) \leq \mathcal{I}(P_{\max}(j))$. In other words, the contribution of $P'(j)$ to $\mathcal{I}(\mathcal{P}^*)$ cannot exceed the contribution of $P_{\max}(j)$ to $\mathcal{I}(\mathcal{P}^*)$. Therefore, $\mathcal{I}(\mathcal{P}^*) \leq \sum_{e_i \in E} \mathcal{I}(P_{\max}(i))$ in this case as well. \square

We now present a lower bound for $\mathcal{I}(\mathcal{P}^*)$.

Lemma 5.3. $\mathcal{I}(\mathcal{P}^*) \geq \sum_{e_i \in E} (\omega(e_i) \times \log 2)$.

Proof. Since \mathcal{P}^* covers all edges in E , a trivial (albeit not necessarily optimal) solution \mathcal{P}' for IP can be constructed by including every edge as a distinct interesting path in the graph, i.e., $\mathcal{P}' = \{e_i | 1 \leq i \leq m\}$. Therefore, $\mathcal{I}(\mathcal{P}^*) \geq \mathcal{I}(\mathcal{P}') = \sum_{e_i \in E} (\omega(e_i) \times \log 2)$ follows from Equation (1). \square

6 Implementation and Testing

We have implemented the core algorithms for long interesting path detection as part of the HYPPO-X repository, which includes our open source implementation of the Mapper framework. The HYPPO-X software repository is open source, and is available at <https://xperthut.github.io/HYPPO-X>. The core computational modules are implemented in C++ and the visualization modules are implemented using the Javascript visualization library, D3 [23].

Currently, the HYPPO-X repository includes implementations for the following algorithms presented in this paper:

1. the optimized exact algorithm for MAX-IP on DAGs described in Section 3.2.1;

2. the greedy iterative heuristic for IP on DAGs described in Algorithm 1; and
3. greedy iterative heuristics for k -IP and ATLEAST- k -IP on DAGs as described in Section 5.1.1.

6.1 Experimental testing

We have been testing and evaluating our implementations on data sets obtained from an ongoing plant phenomics research project. Phenomics is an emerging branch of modern biology that involves the analysis of multiple types of data (genomic, phenotypic, and environmental) acquired using high-throughput technologies. A core goal of plant phenomics research is to understand how different crop genotypes (G) interact with their environments (E) to produce varying performance traits (phenotypes (P)); this goal is often summarized as $G \times E \rightarrow P$ [14].

We present sample results from our analysis of a real world maize data set (accessible from the HYPPO-X repository). This data set consists of phenotypic and environmental measurements for two maize genotypes (abbreviated here for simplicity as A and B), grown in two geographic locations (Nebraska (NE) and Kansas (KS)). The data consists of daily measurements of growth rate alongside multiple environmental variables, over the course of the entire growing season (100 days). In our analysis, each “point” refers to a unique [genotype, location, time] combination. Consequently, the above data set consists of 400 points. Here, “time” was measured as Days After Planting (DAP).

Figure 8 shows an example screenshot of an output generated by running our greedy heuristic for ATLEAST- k -IP on a subset of the above data set containing only genotype B points. For the purpose of this analysis, we used humidity and DAP as two filter functions (i.e., independent variables) and growth rate as our target function (i.e., dependent variable) for crop performance.

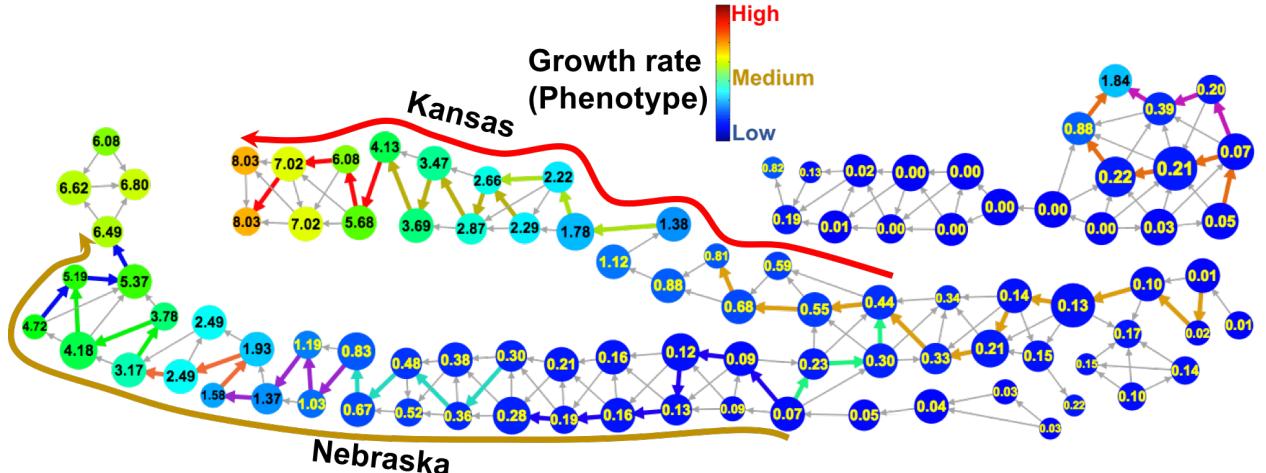


Figure 8: A screenshot of our HYPPO-X tool for identifying interesting paths in Mapper. The results correspond to our experiments with a plant phenomics data set (real world maize crop data). An interactive version of this visualization can be accessed at <https://xperthut.github.io/HYPPO-X>.

Our tool uncovered two long interesting paths—the red-colored path corresponds to a trail of clusters from the Kansas subpopulation, whereas the orange-colored path originates from the Nebraska subpopulation. Note that the location information was *not* used *a priori* by our framework—our path detection was entirely unsupervised! More importantly, the two long interesting paths helped us detect growth rate variations that are potentially tied to specific variations in both humidity and DAP (time), between the crops of these two locations (KS vs NE). We are currently working with our plant science collaborators to better understand the basis for these observed variations, with an aim to postulate hypotheses.

Visualizations of the outputs generated in the above analysis are available on the HYPPO-X repository. More details about this plant phenomics use case can be found in our related application-oriented paper [10].

7 Discussion

We have proposed a general framework for quantifying the significance of features in the Mapper in terms of their interestingness scores. The associated optimization problems MAX-IP, k -IP, and IP constitute a new class of nonlinear longest path problems on directed graphs. We have not characterized the complexity of problem IP. Judging from the fact that k -IP is NP-complete even on DAGs, we suspect IP is NP-complete as well.

Our framework for quantifying interesting paths could be modified to quantify branches in flares as well as holes. When the graph G is a DAG, two interesting k -paths that start and end at the same pair of vertices could be characterized as a $2k$ -hole, i.e., a cycle with $2k$ edges. One could alternatively use persistent homology tools to characterize holes—by identifying “long” generators around them. A 2-way branching in a flare could be identified by two interesting k -paths where one path starts off from a vertex in the middle of the other path. Alternatively, we could generalize the definition of interestingness score for a path in Equation (1) to that of a 2-way branch. Subsequently, we could seek to solve the related optimization problems of identifying the most interesting 2-way flare, or to identify a collection of 2-way flares whose total interestingness score is maximized.

While we distinguished the clustering function g from the filter functions f_i for $i = 1, \dots, h$, this distinction is not critically used in our framework. As indicated in Remark 2.5, one could use g simultaneously as a filter function along with the f_i ’s, and the overall analysis should still carry through. More generally, details of how to implement clustering within the Mapper framework has not received much research attention. In initial work on phenomics [10], we obtained better results when using g alone to cluster within Mapper (rather than clustering using several, or even all, of the variables). It would be interesting to characterize the stability of Mapper to varying settings of clustering employed in its construction. For instance, could we identify a “small” subset of variables for use in clustering within Mapper that is optimal in a suitable sense?

While we have proposed an efficient heuristic for IP on DAGs, we are not able to certify the quality of solution obtained by this method. On the other hand, could we devise approximation algorithms for IP or k -IP? One might have to work under some simplifying assumptions on the distribution of weights $\omega(e)$ or on the structure of the graph G . The simplest case to consider appears to that of IP on a DAG with unit weights on all edges.

We study interestingness of features in a given single Mapper. A natural extension to consider would be to characterize the stability of the highly interesting features. Could we incorporate our interestingness scores into the mathematical machinery recently developed to obtain results on stability and statistical convergence of the 1-D Mapper [2, 3]? Another natural extension to explore would be to define and efficiently identify interesting *surfaces* (or manifolds in general) in higher dimensional mappers.

Acknowledgments: This research is supported by the NSF grant DBI-1661348. Krishnamoorthy thanks Frédéric Meunier for discussion on the complexity of MAX-IP while visiting MSRI.

References

- [1] Muthu Alagappan. From 5 to 13: Redefining the positions in basketball. In *MIT Sloan Sports Analytics Conference*, 2012. <http://www.sloansportsconference.com/content/the-13-nba-positions-using-topology-to-identify-the-different-types-of-players/>.

[2] Mathieu Carrière, Bertrand Michel, and Steve Oudot. Statistical analysis and parameter selection for mapper. 2017. [arXiv:1706.00204](https://arxiv.org/abs/1706.00204).

[3] Mathieu Carrière and Steve Oudot. Structure and stability of the one-dimensional mapper. *Foundations of Computational Mathematics*, Oct 2017. [arXiv:1511.05823](https://arxiv.org/abs/1511.05823). [doi:10.1007/s10208-017-9370-z](https://doi.org/10.1007/s10208-017-9370-z).

[4] Tamal K. Dey, Facundo Mémoli, and Yusu Wang. Multiscale mapper: Topological summarization via codomain covers. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '16, pages 997–1013, Philadelphia, PA, USA, 2016. Society for Industrial and Applied Mathematics. [arXiv:1504.03763](https://arxiv.org/abs/1504.03763).

[5] Jack R. Edmonds. Maximum matching and a polyhedron with 0,1-vertices. *Journal of Research of the National Bureau of Standards Section B*, 69:125–130, 1965.

[6] Harold N. Gabow. Data structures for weighted matching and nearest common ancestors with linking. In *Proceedings of the First Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '90, pages 434–443, Philadelphia, PA, USA, 1990. Society for Industrial and Applied Mathematics. [arXiv:1611.07055](https://arxiv.org/abs/1611.07055).

[7] Geoffrey M. Guisewite and Panos M. Pardalos. Algorithms for the single-source uncapacitated minimum concave-cost network flow problem. *Journal of Global Optimization*, 1(3):245–265, 1991. [doi:10.1007/BF00119934](https://doi.org/10.1007/BF00119934).

[8] Timothy S.C. Hinks, Xiaoying Zhou, Karl J. Staples, Borislav D. Dimitrov, Alexander Manta, Tanya Petrossian, Pek Y. Lum, Caroline G. Smith, Jon A. Ward, Peter H. Howarth, Andrew F. Walls, Stephan D. Gadola, and Ratko Djukanovi. Innate and adaptive T cells in asthmatic patients: Relationship to severity and disease mechanisms. *Journal of Allergy and Clinical Immunology*, 136(2):323–333, 2015. [doi:10.1016/j.jaci.2015.01.014](https://doi.org/10.1016/j.jaci.2015.01.014).

[9] David Houle, Diddahally R Govindaraju, and Stig Omholt. Phenomics: the next challenge. *Nature reviews genetics*, 11(12):855, 2010.

[10] Methun Kamruzzaman, Ananth Kalyanaraman, Bala Krishnamoorthy, and Patrick Schnable. Toward a scalable exploratory framework for complex high-dimensional phenomics data. 2017. [arXiv:1707.04362](https://arxiv.org/abs/1707.04362).

[11] Richard M. Karp. Reducibility Among Combinatorial Problems. In R. E. Miller and J. W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972.

[12] Li Li, Wei-Yi Cheng, Benjamin S. Glicksberg, Omri Gottesman, Ronald Tamler, Rong Chen, Erwin P. Bottinger, and Joel T. Dudley. Identification of type 2 diabetes subgroups through topological analysis of patient similarity. *Science Translational Medicine*, 7(311):311ra174–311ra174, 2015. [doi:10.1126/scitranslmed.aaa9364](https://doi.org/10.1126/scitranslmed.aaa9364).

[13] Pek Y. Lum, Gurjeet Singh, Alan Lehman, Tigran Ishkanov, Mikael. Vejdemo-Johansson, Muthi Alagappan, John G. Carlsson, and Gunnar Carlsson. Extracting insights from the shape of complex data using topology. *Scientific Reports*, 3(1236), 2013. [doi:10.1038/srep01236](https://doi.org/10.1038/srep01236).

[14] Cathie Martin. The plant science decadal vision. *The Plant Cell*, 25(12):4773–4774, 2013.

[15] Yukio Matsumoto. *An Introduction to Morse Theory*. Europe and Central Asia Poverty Reduction and Economic Manag. American Mathematical Society, 2002.

- [16] John W. Milnor. *Morse Theory*. Annals of mathematics studies. Princeton University Press, 1963.
- [17] James R. Munkres. *Elements of Algebraic Topology*. Addison–Wesley Publishing Company, Menlo Park, 1984.
- [18] Monica Nicolau, Arnold J. Levine, and Gunnar Carlsson. Topology based data analysis identifies a subgroup of breast cancers with a unique mutational profile and excellent survival. *Proceedings of the National Academy of Sciences*, 108(17):7265–7270, 2011. [doi:10.1073/pnas.1102826108](https://doi.org/10.1073/pnas.1102826108).
- [19] Jessica L. Nielson, Jesse Paquette, Aiwen W. Liu, Cristian F. Guandique, C. Amy Tovar, Tomoo Inoue, Karen-Amanda Irvine, John C. Gensel, Jennifer Kloke, Tanya C. Petrossian, Pek Y. Lum, Gunnar E. Carlsson, Geoffrey T. Manley, Wise Young, Michael S. Beattie, Jacqueline C. Bresnahan, and Adam R. Ferguson. Topological data analysis for discovery in preclinical spinal cord injury and traumatic brain injury. *Nature Communications*, 6:8581+, October 2015. [doi:10.1038/ncomms9581](https://doi.org/10.1038/ncomms9581).
- [20] Matteo Rucco, Emanuela Merelli, Damir Herman, Devi Ramanan, Tanya Petrossian, Lorenzo Falsetti, Cinzia Nitti, and Aldo Salvi. Using topological data analysis for diagnosis pulmonary embolism. *Journal of Theoretical and Applied Computer Science*, 9:41–55, 2015. <http://www.jtacs.org/archive/2015/1/5>.
- [21] Ghanashyam Sarikonda, Jeremy Pettus, Sonal Phatak, Sowbarnika Sachithanantham, Jacqueline F. Miller, Johnna D. Wesley, Eithon Cadag, Ji Chae, Lakshmi Ganesan, Ronna Mallios, Steve Edelman, Bjoern Peters, and Matthias von Herrath. CD8 T-cell reactivity to islet antigens is unique to type 1 while CD4 T-cell reactivity exists in both type 1 and type 2 diabetes. *Journal of Autoimmunity*, 50(Supplement C):77–82, 2014. [doi:10.1016/j.jaut.2013.12.003](https://doi.org/10.1016/j.jaut.2013.12.003).
- [22] Gurjeet Singh, Facundo Memoli, and Gunnar Carlsson. Topological Methods for the Analysis of High Dimensional Data Sets and 3D Object Recognition. In M. Botsch, R. Pajarola, B. Chen, and M. Zwicker, editors, *Proceedings of the Symposium on Point Based Graphics*, pages 91–100, Prague, Czech Republic, 2007. Eurographics Association. [doi:10.2312/SPBG/SPBG07/091-100](https://doi.org/10.2312/SPBG/SPBG07/091-100).
- [23] Swizec Teller. *Data Visualization with d3.js*. Packt Publishing Ltd, 2013.
- [24] Brenda Y. Torres, Jose Henrique M. Oliveira, Ann Thomas Tate, Poonam Rath, Katherine Cumnock, and David S. Schneider. Tracking resilience to infections by mapping disease space. *PLoS Biol*, 14(4):1–19, 04 2016. [doi:10.1371/journal.pbio.1002436](https://doi.org/10.1371/journal.pbio.1002436).
- [25] George Tsaggouris and Christos Zaroliagis. *Non-additive Shortest Paths*, pages 822–834. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004. [doi:10.1007/978-3-540-30140-0_72](https://doi.org/10.1007/978-3-540-30140-0_72).
- [26] Loring W. Tu. *An Introduction to Manifolds*. Universitext. Springer New York, 2nd edition, 2011.
- [27] Hoang Tuy, Saied Ghannadan, Athanasios Migdalas, and Peter Värbrand. The minimum concave cost network flow problem with fixed numbers of sources and nonlinear arc costs. *Journal of Global Optimization*, 6(2):135–151, Mar 1995. [doi:10.1007/BF01096764](https://doi.org/10.1007/BF01096764).