

Compact Flow Diagrams for State Sequences

Kevin Buchin¹(✉), Maike Buchin², Joachim Gudmundsson³, Michael Horton³,
and Stef Sijben²

¹ Department of Mathematics and Computer Science,
TU Eindhoven, Eindhoven, The Netherlands
`k.a.buchin@tue.nl`

² Department of Mathematics, Ruhr-Universität Bochum, Bochum, Germany
`{Maike.Buchin,Stef.Sijben}@ruhr-uni-bochum.de`

³ School of Information Technologies, The University of Sydney, Sydney, Australia
`{joachim.gudmundsson,michael.horton}@sydney.edu.au`

Abstract. We introduce the concept of compactly representing a large number of state sequences, e.g., sequences of activities, as a flow diagram. We argue that the flow diagram representation gives an intuitive summary that allows the user to detect patterns among large sets of state sequences. Simplified, our aim is to generate a small flow diagram that models the flow of states of all the state sequences given as input. For a small number of state sequences we present efficient algorithms to compute a minimal flow diagram. For a large number of state sequences we show that it is unlikely that efficient algorithms exist. More specifically, the problem is $W[1]$ -hard if the number of state sequences is taken as a parameter. We thus introduce several heuristics for this problem. We argue about the usefulness of the flow diagram by applying the algorithms to two problems in sports analysis. We evaluate the performance of our algorithms on a football data set and generated data.

1 Introduction

Sensors are tracking the activity and movement of an increasing number of objects, generating large data sets in many application domains, such as sports analysis, traffic analysis and behavioural ecology. This leads to the question of how large sets of sequences of activities can be represented compactly. We introduce the concept of representing the “flow” of activities in a compact way and argue that this is helpful to detect patterns in large sets of state sequences.

To describe the problem we start by giving a simple example. Consider three objects (people) and their sequences of states, or activities, during a day. The set of state sequences $\mathcal{T} = \{\tau_1, \tau_2, \tau_3\}$ are shown in Fig. 1(a). As input we are also given a set of criteria $\mathcal{C} = \{C_1, \dots, C_k\}$, as listed in Fig. 1(b). Each criterion is a Boolean function on a single subsequence of states, or a set of subsequences of states. For example, in the given example the criterion $C_1 = \text{“eating”}$ is true for Person 1 at time intervals 7–8 am and 7–9 pm, but false for all other time intervals. Thus, a criterion partitions a sequence of states into subsequences, called *segments*. In each segment the criterion is either true or false. A *segmentation* of

	Person 1	Person 2	Person 3
7-8am	breakfast	gym	breakfast
8-9am	cycle to work	drive to work	cycle to work
9am-5pm	work	work	work
5-7pm	study	dinner	shop
7-9pm	dinner	shop	dinner


(a)

- C_1 : Eating {breakfast,dinner}
- C_2 : Commuting {cycle/drive to work}
- C_3 : Exercising {gym,cycle to work}
- C_4 : Working or studying
- C_5 : Working for at least 4 hours
- C_6 : Shopping
- C_7 : At least 2 people eating simultaneously

(b)

	Person 1	Person 2	Person 3
7-8am	$[C_1, C_7]$	$[C_3]$	$[C_1, C_7]$
8-9am	$[C_2, C_3]$	$[C_2]$	$[C_2, C_3]$
9am-5pm	$[C_4, C_5]$	$[C_4, C_5]$	$[C_4, C_5]$
5-7pm	$[C_4]$	$[C_1]$	$[C_6]$
7-9pm	$[C_1, C_7]$	$[C_6]$	$[C_1, C_7]$

(c)



(d)

Fig. 1. The input is (a) a set $\mathcal{T} = \{\tau_1, \dots, \tau_m\}$ of sequences of states and (b) a set of criteria $\mathcal{C} = \{C_1, \dots, C_k\}$. (c) The criteria partition the states into a segmentation. (d) A valid flow diagram for \mathcal{T} according to \mathcal{C} .

\mathcal{T} is a partition of each sequence in \mathcal{T} into true segments, which is represented by the corresponding sequence of criteria. If a criterion C is true for a set of subsequences, we say they *fulfil* C . Possible segments of \mathcal{T} according to the set \mathcal{C} are shown in Fig. 1(c). The aim is to summarize segmentations of all sequences efficiently; that is, build a flow diagram \mathcal{F} , starting at a start state s and ending at an end state t , with a small number of nodes such that for each sequence of states τ_i , $1 \leq i \leq m$, there exists a segmentation according to \mathcal{C} which appears as an $s-t$ path in \mathcal{F} . A possible flow diagram is shown in Fig. 1(d). This flow diagram for \mathcal{T} according to \mathcal{C} can be validated by going through a segmentation of each object while following a path in \mathcal{F} from s to t . For example, for Person 1 the $s-t$ path $s \rightarrow C_1 \rightarrow C_2 \rightarrow C_4 \rightarrow C_1 \rightarrow t$ is a valid segmentation.

Now we give a formal description of the problem. A *flow diagram* is a node-labelled DAG containing a source node s and sink node t , and where all other nodes are labelled with a criterion. Given a set \mathcal{T} of sequences of states and a set of criteria \mathcal{C} , the goal is to construct a flow diagram with a minimum number of nodes, such that a segmentation of each sequence of states in \mathcal{T} is represented, that is, included as an $s-t$ path, in the flow diagram. Furthermore (when criteria depend on multiple state sequences, e.g. C_7 in Fig. 1) we require that the segmentations represented in the flow diagram are consistent, i.e. can be jointly realized. The *Flow Diagram problem* thus requires the segmentations of each sequence of states and the minimal flow diagram of the segmentations to be computed. It can be stated as:

Problem 1. Flow Diagram (FD)

Instance: A set of sequences of states $\mathcal{T} = \{\tau_1, \dots, \tau_m\}$, each of length at most n , a set of criteria $\mathcal{C} = \{C_1, \dots, C_k\}$ and an integer $\lambda > 2$.

Question: Is there a flow diagram \mathcal{F} with $\leq \lambda$ nodes, such that for each $\tau_i \in \mathcal{T}$, there exists a segmentation according to \mathcal{C} which appears as an s - t path in \mathcal{F} ?

Even the small example above shows that there can be considerable space savings by representing a set of state sequences as a flow diagram. This is not a lossless representation and comes at a cost. The flow diagram represents the sequence of flow between states, however, the information about an individual sequence of states is lost. As we will argue in Sect. 3, paths representing many segments in the obtained flow diagrams show interesting patterns. We will give two examples. First we consider segmenting the morphology of formations of a defensive line of football players during a match (Fig. 4). The obtained flow diagram provides an intuitive summary of these formations. The second example models attacking possessions as state sequences. The summary given by the flow diagram gives intuitive information about differences in attacking tactics.

Properties of Criteria. The efficiency of the algorithms will depend on properties of the criteria on which the segmentations are based. Here we consider four cases: (i) general criteria without restrictions; (ii) monotone decreasing and independent criteria; (iii) monotone decreasing and dependent criteria; and (iv) fixed criteria. To illustrate the properties we will again use the example in Fig. 1.

A criterion C is *monotone decreasing* [8] for a given sequence of states τ that fulfils C , if all subsequences of τ also fulfil C . For example, if C_4 is fulfilled by a sequence τ then any subsequence τ' of τ will also fulfil C_4 . This is in contrast to criterion C_5 which is not monotone decreasing.

A criterion C is *independent* if checking whether a subsequence τ' of a sequence $\tau_i \in \mathcal{T}$ fulfils C can be achieved without reference to any other sequences $\tau_j \in \mathcal{T}, i \neq j$. Conversely, C is *dependent* if checking that a subsequence τ' of τ_i requires reference to other state sequences in \mathcal{T} . In the above example C_4 is an example of an independent criterion while C_7 is a dependent criterion since it requires that at least two objects fulfil the criterion at the same time.

Related Work. To the best of our knowledge compactly representing sequences of states as flow diagrams has not been considered before. The only related work we are aware of comes from the area of trajectory analysis. Spatial trajectories are a special case of state sequences. A spatial trajectory describes the movement of an object through space over time, where the states are location points, which may also include additional information such as heading, speed, and temperature. For a single trajectory a common way to obtain a compact representation is *simplification* [10]. Trajectory simplification asks to determine a subset of the data that represents the trajectory well in terms of the location over time. If the focus is on characteristics other than the location, then *segmentation* [1, 2, 8] is used to partition a trajectory into a small number of subtrajectories, where each subtrajectory is homogeneous with respect to some characteristic. This allows a trajectory to be compactly represented as a sequence of characteristics.

For multiple trajectories other techniques apply. A large set of trajectories might contain very unrelated trajectories, hence *clustering* may be used. Clustering on complete trajectories will not represent information about interesting parts of trajectories; for this clustering on subtrajectories is needed [6, 12]. A set of trajectories that forms different groups over time may be captured by a *grouping structure* [7]. These approaches also focus on location over time.

For the special case of spatial trajectories, a flow diagram can be illustrated by a simple example: trajectories of migrating geese, see [9]. The individual trajectories can be segmented into phases of activities such as directed flight, foraging and stop overs. This results in a flow diagram containing a path for the segmentation of each trajectory. More complex criteria can be imagined that depend on a group of geese, or frequent visits to the same area, resulting in complex state sequences that are hard to analyze without computational tools.

Results, Organization and Hardness. In Sect. 2 we present algorithms for the Flow Diagram problem using criteria with the properties described above. These algorithms only run in polynomial time if the number of state sequences m is constant. Below we observe that this is essentially the best we can hope for by showing that the problem is $W[1]$ -hard.

Theorem 2. *The FD problem is NP-hard. This even holds when only two criteria are used or when the length of every state sequence is 2. Furthermore, for any $0 < c < 1/4$, the FD problem cannot be approximated within factor of $c \log m$ in polynomial time unless $NP \subset DTIME(m^{\text{poly} \log m})$.*

Also for bounded m the running times of our algorithms is rather high. Again, we can show that there are good reasons for this.

Theorem 3. *The FD problem parameterized in the number of state sequences is $W[1]$ -hard even when the number of criteria is constant.*

Both theorems are proved in the longer version of this paper [5]. Unless $W[1] = FPT$, this rules out the existence of algorithms with time complexity of $O(f(m) \cdot (nk)^c)$ for some constant c and any computable function $f(m)$, where m, n and k are the number of state sequences, the length of the state sequences and the number of criteria, respectively. To obtain flow diagrams for larger groups of state sequences we propose two heuristics for the problem in Sect. 2. We experimentally evaluate the algorithms and heuristics in Sect. 3.

2 Algorithms

In this section, we present algorithms that compute a smallest flow diagram representing a set of m state sequences of length n for a set of k criteria. First, we present an algorithm for the general case, followed by a more efficient algorithm for the case of monotone increasing and independent criteria, and then two heuristic algorithms. The algorithm for monotone increasing and dependent criteria, and the proofs omitted in this section are in the extended version of this paper [5].

2.1 General Criteria

Next, we present a dynamic programming algorithm for finding a smallest flow diagram. Recall that a node v in the flow diagram represents a criterion C_j that is fulfilled by a contiguous segment in some of the state sequences. Let $\tau[i, j]$, $i \leq j$, denote the subsequence of τ starting at the i th state of τ and ending at the j th state, where $\tau[i, i]$ is the empty sequence. Construct an $(n + 1)^m$ grid of vertices, where a vertex with coordinates (x_1, \dots, x_m) , $0 \leq x_1, \dots, x_m \leq n$, represents $(\tau_1[0, x_1], \dots, \tau_m[0, x_m])$. Construct a *prefix graph* G as follows:

There is an edge between two vertices $v = (x_1, \dots, x_m)$ and $v' = (x'_1, \dots, x'_m)$, labeled by some criterion C_j , if and only if, for every i , $1 \leq i \leq m$, one of the following two conditions is fulfilled: (1) $x_i = x'_i$, or (2) all remaining $\tau_i[x_i + 1, x'_i]$ jointly fulfil C_j . Consider the edge between $(x_1, x_2) = (1, 0)$ and $(x'_1, x'_2) = (1, 1)$ in Fig. 2(b). Here $x_1 = x'_1$ and $\tau_2[x_2 + 1, x'_2]$ fulfils C_2 .

Finally, define v_s to be the vertex in G with coordinates $(0, \dots, 0)$ and add an additional vertex v_t outside the grid, which has an incoming edge from (n, \dots, n) . This completes the construction of the prefix graph G .




Fig. 2. (a) A segmentation of $\mathcal{T} = \{\tau_1, \tau_2\}$ according to $\mathcal{C} = \{C_1, C_2, C_3\}$. (b) The prefix graph G of the segmentation, omitting all but four of the edges. (c) The resulting flow diagram generated from the highlighted path in the prefix graph.

Now, a path in G from v_s to a vertex v represents a valid segmentation of some prefix of each state sequence, and defines a flow diagram that describes these segmentations in the following way: the empty path represents the flow diagram consisting only of the start node s . Every edge of the path adds one new node to the flow diagram, labeled by the criterion that the segments fulfil. Additionally, for each node the flow diagram contains an edge from every node representing a previous segment, or from s if the node is the first in a segmentation. For a path leading from v_s to v_t , the target node t is added to the flow diagram, together with its incoming edges. This ensures that the flow diagram represents valid segmentations and that each node represents at least one segment. An example of this construction is shown in Fig. 2.

Hence the length of a path (where length is the number of edges on the path) equals the number of nodes of the corresponding flow diagram, excluding s and t . Thus, we find an optimal flow diagram by finding a shortest v_s - v_t path in G .

Lemma 4. *A smallest flow diagram for a given set of state sequences is represented by a shortest v_s - v_t path in G .*

Recall that G has $(n+1)^m$ vertices. Each vertex has $O(k(n+1)^m)$ outgoing edges, thus, G has $O(k(n+1)^{2m})$ edges in total. To decide if an edge is present in G , check if the nonempty segments the edge represents fulfil the criterion. Thus, we need to perform $O(k(n+1)^{2m})$ of these checks. There are m segments of length at most n , and we assume the cost for checking this is $T(m, n)$. Thus, the cost of constructing G is $O(k(n+1)^{2m} \cdot T(m, n))$, and finding the shortest path requires $O(k(n+1)^{2m})$ time.

Theorem 5. *The algorithm described above computes a smallest flow diagram for a set of m state sequences, each of length at most n , and k criteria in $O((n+1)^{2m}k \cdot T(m, n))$ time, where $T(m, n)$ is the time required to check if a set of m subsequences of length at most n fulfils a criterion.*

2.2 Monotone Decreasing and Independent Criteria

If all criteria are decreasing monotone and independent, we can use ideas similar to those presented in [8] to avoid constructing the full graph. From a given vertex with coordinates (x_1, \dots, x_m) , we can greedily move as far as possible along the sequences, since the monotonicity guarantees that this never leads to a solution that is worse than one that represents shorter segments. For a given criterion C_j , we can compute for each τ_i independently the maximum x'_i such that $\tau_i[x_i+1, x'_i]$ fulfils C_j . This produces coordinates (x'_1, \dots, x'_m) for a new vertex, which is the optimal next vertex using C_j . By considering all criteria we obtain k new vertices. However, unlike the case with a single state sequence, there is not necessarily one vertex that is better than all others (i.e. largest ending position), since there is no total order on the vertices. Instead, we consider all vertices that are not dominated by another vertex, where a vertex p dominates a vertex p' if each coordinate of p is at least as large as the corresponding coordinate of p' , and at least one of p 's coordinates is larger.

Let V_i be the set of vertices of G that are reachable from v_s in exactly i steps, and define $M(V) := \{v \in V \mid \text{no vertex } u \in V \text{ dominates } v\}$ to be the set of *maximal vertices* of a vertex set V . Then a shortest v_s - v_t path through G can be computed by iteratively computing $M(V_i)$ for increasing i , until a value of i is found for which $v_t \in M(V_i)$. Observe that $|M(V)| = O((n+1)^{m-1}k)$ for any set V of vertices in the graph. Also note that $V_0 = M(V_0) = v_s$.

Lemma 6. *For each $i \in \{1, \dots, \ell - 1\}$, every vertex in $M(V_i)$ is reachable in one step from a vertex in $M(V_{i-1})$. Here, ℓ is the distance from v_s to v_t .*

$M(V_i)$ is computed by computing the farthest reachable vertex for each $v \in M(V_{i-1})$ and criterion, thus yielding a set D_i of $O((n+1)^{m-1}k)$ vertices. This set contains $M(V_i)$ by Lemma 6, so we now need to remove all vertices that are dominated by some other vertex in the set to obtain $M(V_i)$.

We find $M(V_i)$ using a copy of G . Each vertex may be marked as being in D_i or dominated by a vertex in D_i . We process the vertices of D_i in arbitrary

order. For a vertex v , if it is not yet marked, we mark it as being in D_i . When a vertex is newly marked, we mark its $\leq m$ immediate neighbours dominated by it as being dominated. After processing all vertices, the grid is scanned for the vertices still marked as being in D_i . These vertices are exactly $M(V_i)$.

When computing $M(V_i)$, $O((n+1)^{m-1}k)$ vertices need to be considered, and the maximum distance from v_s to v_t is $m(n+1)$, so the algorithm considers $O(mk(n+1)^m)$ vertices. We improve this bound by a factor m using the following:

Lemma 7. *The total size of all D_i , for $0 \leq i \leq \ell - 1$, is $O(k(n+1)^m)$.*

Using this result, we compute all $M(V_i)$ in $O((k+m)(n+1)^m)$ time, since $O(k(n+1)^m)$ vertices are marked directly, and each of the $(n+1)^m$ vertices is checked at most m times when a direct successor is marked. One copy of the grid can be reused for each $M(V_i)$, since each vertex of D_{i+1} dominates at least one vertex of $M(V_i)$ and is thus not yet marked while processing D_j for any $j \leq i$.

Since the criteria are independent, the farthest reachable point for a given starting point and criterion can be precomputed for each state sequence separately. Using the monotonicity we can traverse each state sequence once per criterion and thus need to test only $O(nmk)$ times whether a subsequence fulfils a criterion.

Theorem 8. *The algorithm described above computes a smallest flow diagram for m state sequences of length n with k independent and monotone decreasing criteria in $O(mnk \cdot T(1, n) + (k+m)(n+1)^m)$ time, where $T(1, n)$ is the time required to check if a subsequence of length at most n fulfils a criterion.*

2.3 Heuristics

The hardness results presented in the introduction indicate that it is unlikely that the performance of the algorithms will be acceptable in practical situations, except for very small inputs. As such, we investigated heuristics that may produce usable results that can be computed in reasonable time.

We consider heuristics for monotone decreasing and independent criteria. These are based on the observation that by limiting V_i , the vertices that are reachable from v_s in i steps, to a fixed size, the complexity of the algorithm can be controlled. Given that every path in a prefix graph represents a valid flow diagram, any path chosen in the prefix graph will be valid, though not necessarily optimal. In the worst case, a vertex that advances along a single state sequence a single time-step (i.e. advancing only one state) will be selected, and for each vertex, all k criteria must be evaluated, so $O(kmn)$ vertices may be processed by the algorithm. We consider two strategies for selecting the vertices in V_i to retain:

(1) For each vertex in V_i , determine the number of state sequences that are advanced in step i and retain the top q vertices [*sequence heuristic*].

(2) For each vertex in V_i , determine the number of time-steps that are advanced in all state sequences in step i and retain the top q vertices [*time-step heuristic*].

In our experiments we use $q = 1$ since any larger value would immediately give an exponential worst-case running time.

3 Experiments

The objectives of the experiments were twofold: to determine whether compact and useful flow diagrams could be produced in real application scenarios; and to empirically investigate the performance of the algorithms on inputs of varying sizes. We implemented the algorithms described in Sect. 2 using the Python programming language. For the first objective, we considered the application of flow diagrams to practical problems in football analysis in order to evaluate their usefulness. For the second objective, the algorithms were run on generated datasets of varying sizes to investigate the impact of different parameterisations on the computation time required to produce the flow diagram and the complexity of the flow diagram produced.

3.1 Tactical Analysis in Football

Sports teams will apply tactics to improve their performance, and computational methods to detect, analyse and represent tactics have been the subject of several recent research efforts [4, 11, 14, 16–18]. Two manifestations of team tactics are in the persistent and repeated occurrence of spatial formations of players, and in *plays* — a coordinated sequence of actions by players. We posited that flow diagrams would be a useful tool for compactly representing both these manifestations, and we describe the approaches used in this section.

The input for the experiments is a database containing player trajectory and match event data from four home matches of the Arsenal Football Club from the 2007/08 season, provided by Prozone Sports Limited [15]. For each player and match, there is a trajectory comprising a sequence of timestamped location points in the plane, sampled at 10 Hz and accurate to 10 cm. The origin of the coordinate system coincides with the centre point of the football pitch and the longer side of the pitch is parallel to the x -axis — i.e. the pitch is oriented so the goals are to the left and right. In addition, for each match, there is a log of all the match events, comprising the type, time-stamp and location of each event.

Defensive Formations. The spatial formations of players in football matches are known to characterize a team’s tactics [3], and a compact representation of how formations change over time would be a useful tool for analysis. We investigated whether a flow diagram could provide such a compact representation of the defensive formation of a team, specifically to show how the formation evolves during a phase of play. In our match database, all the teams use a formation of four defensive players who orient themselves in line across the pitch. Broadly speaking, the ideal is for the formation to be “flat”, i.e. the players are positioned in a line parallel to the y -axis. However the defenders will react to changes circumstances, for example in response to opposition attacks, possibly

causing the formation to deform. We constructed the following flow diagram to analyse the defensive formations used in the football matches in our database.

For each match in the database, the trajectories of the four defensive players were re-sampled at one-second intervals to extract the point-locations of the four defenders. The samples were partitioned into sequences $\mathcal{T} = \{\tau_1, \dots, \tau_m\}$ corresponding to phases such that a single team was in possession of the ball, and where the phase began with a goal kick event, or the goalkeeper kicks or throws the ball from hand. Let $\tau_i[j]$ be the j -th state in the i -th state sequence. Each $\tau_i[j] = (p_1, p_2, p_3, p_4)$, where p_i is the location of a player in the plane, such that the locations are ordered by their y -coordinate: $y(p_i) \leq y(p_{i+1}) : i \in \{1, 2, 3\}$.

The criteria used to summarise the formations were derived from those presented by Kim et al. [13]. The angles between pairs of adjacent players (along the defensive line) were used to compute the formation criteria, see Fig. 3. The scheme in Kim et al. was extended to allow multiple criteria to be applied where the angle between pairs of players is close to 10° . The reason for this was to facilitate compact results by allowing for smoothing of small variations in contiguous time-steps.

The criteria \mathcal{C} applied to each state is a triple (x_1, x_2, x_3) , computed as follows. Given two player positions p and q as points in the plane such that $y(p) \leq y(q)$, let p' be an arbitrary point on the interior of the half-line from p in the direction of the positive y -axis, and let $\angle p'pq$ be the angle induced by these points, and thus denotes the angle between the two player's positions relative to the goal-line. Let $R(-1) = [-90^\circ, -5^\circ]$, $R(0) = (-15^\circ, +15^\circ)$, and $R(1) = (+5^\circ, +90^\circ]$ be three angular ranges. Thus, $\mathcal{C} = \{(x_1, x_2, x_3) : x_1, x_2, x_3 \in \{-1, 0, 1\}\}$ is the set of available criteria.

Each state sequence $\tau_i \in \mathcal{T}$ is segmented according to the criteria set \mathcal{C} . A given state $\tau_i[j] = (p_1, p_2, p_3, p_4)$ may satisfy the criteria (and thus have the formation) (x_1, x_2, x_3) if $\angle p'_i p_i p_{i+1} \in R(x_i)$ for all $i \in \{1, 2, 3\}$.

The criteria are monotone decreasing and independent, and we ran the corresponding algorithm using randomly selected sets of the state sequences as input. The size m of the input was increased until the running time exceeded a threshold of 6 h. The algorithm successfully processed up to $m = 12$ state sequences, having a total of 112 assigned segments. The resulting flow diagram, Fig. 4, has a total complexity of 12 nodes and 27 edges.

We believe that the flow diagram provides an intuitive summary of the defensive formation, and several observations are apparent. There appears to be a preference amongst the teams for the right-back to position himself in advance of the right centre-half (i.e. the third component of the triple is $+1$). Furthermore, the $(0, 0, 0)$ triple, corresponding to a ‘‘flat back four’’ is not present in the diagram. This is typically considered the ideal formation for teams that utilise the offside trap, and thus may suggest that the defences here are not employing this tactic. These observations were apparent to the authors as laymen, and we would expect that a domain expert would be able to extract further useful insights from the flow diagrams.

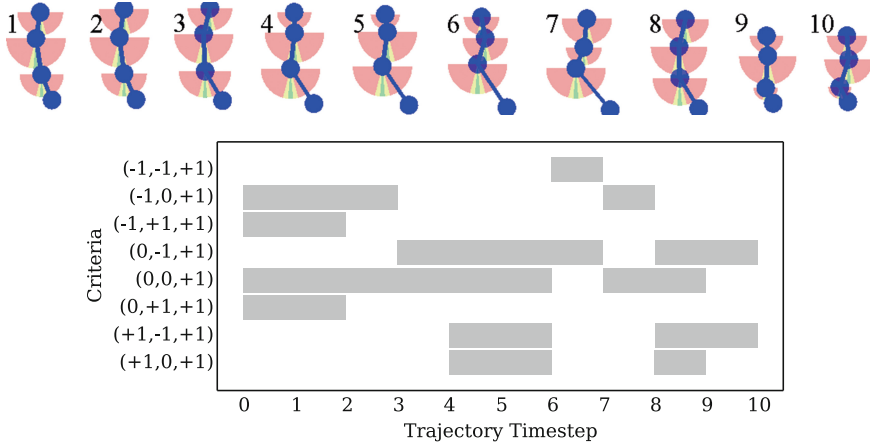


Fig. 3. Segmentation of a single state sequence τ_i . The formation state sequence is used to compute the segmentation representation, where segments corresponding to criteria span the state sequence (*bottom*). The representation of this state sequence in the movement flow diagram is shaded in Fig. 4.




Fig. 4. Flow diagram for formation morphologies of twelve defensive possessions. The shaded nodes are the segmentation of the state sequence in Fig. 3.

Attacking Plays. In this second experiment, we used a different formulation to produce flow diagrams to summarise phases of attack. During a match, the team in possession of the ball regularly attempts to reach a position where they can take a shot at goal. Teams will typically use a variety of tactics to achieve such a position, e.g. teams can vary the intensity of an attack by pushing forward, moving laterally, making long passes, or retreating and regrouping. We modelled attacking possessions as state sequences, segmented according to criteria representing the attacking intensity and tactics employed, and computed flow diagrams for the possessions. In particular, we were interested in determining whether differences in tactics employed by teams when playing at home or away [4] are apparent in the flow diagrams.

We focus on *ball events*, where a player touches the ball, e.g. passes, touches, dribbles, headers, and shots at goal. The event sequence for each match was partitioned into sequences $\mathcal{T} = \{\tau_1, \dots, \tau_m\}$ such that each τ_i is an event sequence

where a single team was in possession, and \mathcal{T} includes only the sequences that end with a shot at goal. Let $\tau_i[j]$ be a tuple (p, t, e) where p is the location in the plane where an event of type $e \in \{touch, pass, dribble, header, shot, clearance\}$ occurred at time t . We are interested in the movement of the ball between an event state $\tau_i[j]$ and the next event state $\tau_i[j+1]$, in particular, let $d_x(\tau_i[j])$ (resp. $d_y(\tau_i[j])$) be the distance in the x -direction (resp. y -direction) between state $\tau_i[j]$ and the next state. Similarly, let $v_x(\tau_i[j])$ (resp. $v_y(\tau_i[j])$) be the velocity of the ball in the x -direction (resp. y -direction) between $\tau_i[j]$ and its successor state. Let $\angle\tau_i[j]$ be the angle defined by the location of $\tau_i[j]$, $\tau_i[j+1]$ and a point on the interior of the half-line from the location of $\tau_i[j]$ in the positive y -direction.

Criteria were defined to characterise the movement of the ball — relative to the goal the team is attacking — between event states in the possession sequence. The criteria $\mathcal{C} = \{C_1, \dots, C_8\}$ were defined as follows.

- C_1 : *Backward movement (BM)*: $v_x(\tau_i[j]) < 1$ — a sub-sequence of passes or touches that move in a defensive direction.
- C_2 : *Lateral movement (LM)*: $-5 < v_x(\tau_i[j]) < 5$ — passes or touches that move in a lateral direction.
- C_3 : *Forward movement (FM)*: $-1 < v_x(\tau_i[j]) < 12$ — passes or touches that move in an attacking direction, at a velocity in the range achievable by a player sprinting, i.e. approximately 12 m/s.
- C_4 : *Fast forward movement (FFM)*: $8 < v_x(\tau_i[j])$ — passes or touches moving in an attacking direction at a velocity generally in excess of maximum player velocity.
- C_5 : *Long ball (LB)*: $30 < d_x(\tau_i[j])$ — a single pass travelling 30 m in the attacking direction.
- C_6 : *Cross-field ball (CFB)*: $20 < d_y(\tau_i[j]) \wedge \angle\tau_i[j] \in [-10, 10] \cup [170, 190]$ — a single pass travelling 20 m in the cross-field direction with an angle within 10° of the y -axis.
- C_7 : *Shot resulting in goal (SG)*: a successful shot resulting in a goal.
- C_8 : *Shot not resulting in goal (SNG)*: a shot that does not produce a goal.

For a football analyst, the first four criteria are simple movements, and are not particularly interesting. The last four events are significant: the long ball and cross-field ball change the locus of attack; and the shot criteria represent the objective of an attack.

The possession state sequences for the home and visiting teams were segmented according to the criteria and the time-step heuristic algorithm was used to compute the flow diagrams. The home-team input consisted of 66 sequences covered by a total of 866 segments, and resulted in a flow diagram with 25 nodes and 65 edges, see Fig. 5. Similarly, the visiting-team input consisted of 39 state sequences covered by 358 segments and the output flow diagram complexity was 22 nodes and 47 edges, as shown in Fig. 6.

At first glance, the differences between these flow diagrams may be difficult to appreciate, however closer inspection reveals several interesting observations. The s - t paths in the home-team flow diagram tend to be longer than those in the visiting team's, suggesting that the home team tends to retain possession of




Fig. 5. Flow diagrams produced for home team. The edge weights are the number of possessions that span the edge, and the nodes with grey background are event types that are significant.

the ball for longer, and varies the intensity of attack more often. Moreover, the nodes for cross-field passes and long-ball passes tend to occur earlier in the $s-t$ paths in the visiting team’s flow diagram. These are both useful tactics as they alter the locus of attack, however they also carry a higher risk. This suggests that the home team is more confident in its ability to maintain possession for long attack possessions, and will only resort to such risky tactics later in a possession. Furthermore, the tactics used by the team in possession are also impacted by the defensive tactics. As Bialkowski et al. [4] found, visiting teams tend to set their defence deeper, i.e. closer to the goal they are defending. When the visiting team is in possession, there is thus likely to be more space behind the home team’s defensive line, and the long ball may appear to be a more appealing tactic. The observations made from these are consistent with our basic understanding of football tactics, and suggest that the flow diagrams are interpretable in this application domain.




Fig. 6. Flow diagrams produced for visiting team. The edge weights are the number of possessions that span the edge, and the nodes with grey background are event types that are significant.

3.2 Performance Testing

In the second experiment, we used a generator that outputs synthetic state sequences and segmentations, and tested the performance of the algorithms on inputs of varying sizes.

The segmentations were generated using Markov-Chain Monte-Carlo sampling. Nodes representing the criteria set of size k were arranged in a ring and a Markov chain constructed, such that each node had a transition probability of 0.7 to remain at the node, 0.1 to move to the adjacent node, and 0.05 to move to the node two places away. Segmentations were computed by sampling the Markov chain starting at a random node. Thus, simulated datasets of arbitrary size m , state sequence length n , criteria set size k were generated.

We performed two tests on the generated segmentations. In the first, experiments were run on the four algorithms described in Sect. 2 with varying configurations of m , n and k to investigate the impact of input size on the algorithm’s performance. The evaluation metric used was the CPU time required to generate the flow diagram for the input. In the second test, we compared the total complexity of the output flow diagram produced by the two heuristic algorithms with




Fig. 7. Runtime statistics for generating flow diagram (*top*), and total complexity of flow diagrams produced (*bottom*). Default values of $m = 4$, $n = 4$ and $k = 10$ were used. The data points are the mean value and the error bars delimit the range of values over the five trials run for each input size.

the baseline complexity of the flow diagram produced by the exact algorithm for monotone increasing and independent criteria.

We repeated each experiment five times with different input sequences for each trial, and the results presented are the mean values of the metrics over the trials. Limits were set such that the process was terminated if the CPU time exceeded 1 h, or the memory required exceeded 8 GB.

The results of the first test showed empirically that the exact algorithms have time and storage complexity consistent with the theoretical worst-case bounds, Fig. 7 (*top*). The heuristic algorithms were subsequently run against larger test data sets to examine the practical limits of the input sizes, and were able to process larger input — for example, an input of $k = 128$, $m = 32$ and $n = 1024$ was tractable — the trade-off is that the resulting flow diagrams were suboptimal, though correct, in terms of their total complexity.

For the second test, we investigated the complexity of the flow diagram induced by inputs of varying parameterisations when using the heuristic algorithms. The objective was to examine how close the complexity was to the

optimal complexity produced using an exact algorithm. The inputs exhibited monotone decreasing and independent criteria, and thus the corresponding algorithm was used to produce the baseline. Figure 7 (*bottom*) summarises the results for varying input parameterisations. The complexity of the flow diagrams produced by the two heuristic algorithms are broadly similar, and increase at worst linearly as the input size increases. Moreover, while the complexity is not optimal it appears to remain within a constant factor of the optimal, suggesting that the heuristic algorithms could produce usable flow diagrams for inputs where the exact algorithms are not tractable.

4 Concluding Remarks

We introduced flow diagrams as a compact representation of a large number of state sequences. We argued that this representation gives an intuitive summary allowing the user to detect patterns among large sets of state sequences, and gave several algorithms depending on the properties of the segmentation criteria. These algorithms only run in polynomial time if the number of state sequences m is constant, which is the best we can hope for given the problem is $W[1]$ -hard. As a result we considered two heuristics capable of processing large data sets in reasonable time, however we were unable to give an approximation bound. We tested the algorithms experimentally to assess the utility of the flow diagram representation in a sports analysis context, and also analysed the performance of the algorithms of inputs of varying parameterisations.

References

1. Alewijnse, S.P.A., Buchin, K., Buchin, M., Kölzsch, A., Kruckenberg, H., Westenberg, M.: A framework for trajectory segmentation by stable criteria. In: Proceedings of 22nd ACM SIGSPATIAL/GIS, pp. 351–360. ACM (2014)
2. Aronov, B., Driemel, A., van Kreveld, M.J., Löffler, M., Staals, F.: Segmentation of trajectories for non-monotone criteria. In: Proceedings of 24th ACM-SIAM SODA, pp. 1897–1911 (2013)
3. Bialkowski, A., Lucey, P., Carr, G.P.K., Yue, Y., Sridharan, S., Matthews, I.: Identifying team style in soccer using formations learned from spatiotemporal tracking data. In: ICDM Workshops, pp. 9–14. IEEE (2014)
4. Bialkowski, A., Lucey, P., Carr, P., Yue, Y., Matthews, I.: Win at home and draw away: automatic formation analysis highlighting the differences in home and away team behaviors. In: Proceedings of 8th Annual MIT Sloan Sports Analytics Conference (2014)
5. Buchin, K., Buchin, M., Gudmundsson, J., Horton, M., Sijben, S.: Compact flow diagrams for state sequences. CoRR, abs/1602.05622 (2016)
6. Buchin, K., Buchin, M., Gudmundsson, J., Löffler, M., Luo, J.: Detecting commuting patterns by clustering subtrajectories. *Int. J. Comput. Geom. Appl.* **21**(3), 253–282 (2011)
7. Buchin, K., Buchin, M., van Kreveld, M., Speckmann, B., Staals, F.: Trajectory grouping structure. In: Dehne, F., Solis-Oba, R., Sack, J.-R. (eds.) WADS 2013. LNCS, vol. 8037, pp. 219–230. Springer, Heidelberg (2013)

8. Buchin, M., Driemel, A., van Kreveld, M., Sacristan, V.: Segmenting trajectories: a framework and algorithms using spatiotemporal criteria. *J. spat. inf. sci.* **3**, 33–63 (2011)
9. Buchin, M., Kruckenberg, H., Kölzsch, A.: Segmenting trajectories based on movement states. In: *Proceedings of 15th SDH*, pp. 15–25. Springer (2012)
10. Cao, H., Wolfson, O., Trajcevski, G.: Spatio-temporal data reduction with deterministic error bounds. *VLDB J.* **15**(3), 211–228 (2006)
11. Gudmundsson, J., Wolle, T.: Football analysis using spatio-temporal tools. *Comput. Environ. Urban Syst.* **47**, 16–27 (2014)
12. Han, C.-S., Jia, S.-X., Zhang, L., Shu, C.-C.: Sub-trajectory clustering algorithm based on speed restriction. *Comput. Eng.* **37**(7), 219–221 (2011)
13. Kim, H.-C., Kwon, O., Li, K.-J.: Spatial and spatiotemporal analysis of soccer. In: *Proceedings of 19th ACM SIGSPATIAL/GIS*, pp. 385–388. ACM (2011)
14. Lucey, P., Bialkowski, A., Carr, G.P.K., Morgan, S., Matthews, I., Sheikh, Y.: Representing and discovering adversarial team behaviors using player roles. In: *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2013)*, Portland, pp. 2706–2713. IEEE, June 2013
15. Prozone Sports Ltd: Prozone Sports - Our technology (2015). <http://prozonesports.stats.com/about/technology/>
16. Van Haaren, J., Dzyuba, V., Hannosset, S., Davis, J.: Automatically discovering offensive patterns in soccer match data. In: Fromont, E., De Bie, T., van Leeuwen, M. (eds.) *IDA 2015. LNCS*, vol. 9385, pp. 286–297. Springer, Heidelberg (2015). doi:10.1007/978-3-319-24465-5_25
17. Wang, Q., Zhu, H., Hu, W., Shen, Z., Yao, Y.: Discerning tactical patterns for professional soccer teams. In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD 2015*, Sydney, pp. 2197–2206. ACM Press, August 2015
18. Wei, X., Sha, L., Lucey, P., Morgan, S., Sridharan, S.: Large-scale analysis of formations in soccer. In: *2013 International Conference on Digital Image Computing: Techniques and Applications (DICTA)*, Hobart, pp. 1–8. IEEE, November 2013