Social Network Analysis and Mining (2019) 9:58
https://doi.org/10.1007/s13278-019-0600-z

ORIGINAL ARTICLE

Mining and modeling complex leadership–followership dynamics of movement data

Chainarong Amornbunchornvej[1] · Tanya Y. Berger-Wolf[2]

Received: 30 November 2018 / Revised: 6 June 2019 / Accepted: 16 September 2019 / Published online: 3 October 2019

Abstract
Leadership and followership are essential parts of collective decision and organization in social animals, including humans. In nature, relationships of leaders and followers are dynamic and vary with context or temporal factors. Understanding dynamics of leadership and followership, such as how leaders and followers change, emerge, or converge, allows scientists to gain more insight into group decision-making and collective behavior in general. However, given only data of individual activities, it is challenging to infer the dynamics of leaders and followers. In this paper, we focus on mining and modeling frequent patterns of leading and following. We formalize new computational problems and propose a framework that can be used to address several questions regarding group movement. We use the leadership inference framework, mFLICA, to infer the time series of leaders and their factions from movement datasets and then propose an approach to mine and model frequent patterns of both leadership and followership dynamics. We evaluate our framework performance by using several simulated datasets, as well as the real-world dataset of baboon movement to demonstrate the applications of our framework. These are novel computational problems and, to the best of our knowledge, there are no existing comparable methods to address them. Thus, we modify and extend an existing leadership inference framework to provide a non-trivial baseline for comparison. Our framework performs better than this baseline in all datasets. Our framework opens the opportunities for scientists to generate testable scientific hypotheses about the dynamics of leadership in movement data.

Keywords Leadership · Followership · Coordination · Time series · Collective behavior

1 Introduction

Leadership is a process that leaders influence followers' actions in order to achieve the collective goal (Hogg 2001; Glowacki and von Rueden 2015). Leadership is an essential part that fosters success of coordinated behaviors in social species (Glowacki and von Rueden 2015; Couzin et al. 2005; Hogg 2001), such as foraging, migration, territorial defense, and so on. In most species, leadership is not permanent but may change with context (the one who leads the group to food or water may be different from the one who leads the flight from a predator) or other social circumstances (two rivaling subgroups may come to a joint decision and merge under single leadership or, vice versa, a group may split to explore several directions). Understanding dynamics of leadership, such as how leaders change, emerge, or converge, allows scientists to gain more insight into group decision-making and collective behavior in general. In this paper, we focus on mining and modeling frequent patterns of leadership dynamics.

One of the intuitive definitions of leadership that is commonly found in nature is the initiation of coordinated activities (Krause et al. 2000; Smith et al. 2015; Stueckle and Zinner 2008). In the context of movement, leaders are initiators who initiate coordinated movement that everyone follows (Amornbunchornvej and Berger-Wolf 2018a). There are several works have been developed to infer leaders from time series of movement data, such as FLOCK method (Andersson et al. 2008), LPD framework (Kjargaard et al. 2013), and methods based on a dynamic

✉ Chainarong Amornbunchornvej
chainarong.amo@nectec.or.th

Tanya Y. Berger-Wolf
tanyabw@uic.edu

[1] National Electronics and Computer Technology Center (NECTEC), Pathum Thani, Thailand

[2] Department of Computer Science, University of Illinois at Chicago, Chicago, IL, USA

Table 1 Comparison of frameworks that can detect clusters in time series

Properties\frameworks	MONIC	TRACLUS, TCMM	mFLICA
Cluster types	Static clusters	Temporal clusters	Factions
Members of clusters	Data points	Time series	Time series
Tracking clusters dynamics	Yes	No	Yes
Following relations	No	No	Yes

Static clusters are clusters that are defined over data points in each time step, while temporal clusters are defined over segments of time series. Factions are temporal clusters that all members follow the same leader. Tracking clusters dynamics implies that a framework can track evolution of clusters, such as merging or splitting of clusters over time. The following relation property implies that a framework can give a relation of who follows whom for all pairs of members within a cluster

following network concept (Amornbunchornvej et al. 2018; Amornbunchornvej and Berger-Wolf 2018a).

Nevertheless, the challenges in the field still remain regarding how to infer and model *the dynamics of the frequent patterns* of leadership events. For example, suppose i and j lead separate subgroups, how often do the two groups merge to a larger group lead by k? How likely is it that the group lead by i will split into more than three subgroups?

However, only the state-of-the-art approach, mFLICA (Amornbunchornvej and Berger-Wolf 2018a), is capable of inferring dynamics of leadership—i.e., emergence, convergence, or a change of leaders—during coordinated movement. mFLICA detects clusters (factions) based on the concept of following relations. In mFLCIA, the time series from the same faction must follow the same leader.

There are many works focusing on inferring dynamics of groups or clusters (Li et al. 2010; Lee et al. 2007; Spiliopoulou et al. 2006). The work by Spiliopoulou et al. (2006) proposed a framework named ’MONIC’ to track various types of clusters transition in time series, such as expanding, splitting, and merging. However, MONIC infers clusters based on time points without considering following relations among time series to detect clusters. The work by Li et al. (2010), Lee et al. (2007) proposed frameworks (TRACLUS and TCMM) to detect temporal clusters from segments of time series. Nevertheless, the temporal clusters are measured based on trajectory similarity without the following relation property. Hence, MONIC, TRACLUS, and TCMM frameworks cannot be used to detect factions of time series, which implies that they cannot detect leadership and followership dynamics. Table 1 summarizes the comparison of these frameworks.

In this paper, we focus on mining frequent patterns of leadership dynamics that requires our framework to identify both groups and leaders of those groups that change over time. Moreover, since the groups following a leader during coordination have a special structure of following relations among members, the standard clustering methods cannot be used in this case.

> Mining Patterns of Leadership Dynamics: Given time series of individual activities, the goal is to mine and model frequent patterns of leadership dynamics, including emergence of multiple leaders, convergence of multiple leaders to a single one, or change of a leader.

1.1 Previous contributions: leadership dynamics

To address these computational questions, in the previous paper in Amornbunchornvej and Berger-Wolf (2018b), we formalize the problem of Mining Patterns of Leadership Dynamics, as well as propose a framework, which is the extension of mFLICA (Amornbunchornvej and Berger-Wolf 2018a), as a solution to this problem. We adapt the traditional framework of frequent pattern mining (Agrawal et al. 1993; Han et al. 2007; Aggarwal and Han 2014) and the Hidden Markov Model (HMM) approach (Rabiner 1989) to model the dynamics of frequent patterns of leadership. Our framework is capable of:

- *Mining and modeling frequent patterns of leadership dynamics* inferring the transition diagram of frequent dynamics of complex leadership events, such as ’a single group lead by k splits into two groups lead by i and j’. In addition, we infer the probabilities of the transitions between such two events in the diagram.
- *Evaluating the significance of leadership-event order* we propose a null model of the dynamics of leadership events and perform hypothesis testing to compare frequent pattern model of leadership dynamics inferred from the given input to our proposed null model.
- *Mining sequence patterns of leadership dynamics* finding support values for the leadership dynamics sequences from time series of movement data.
- *Evaluating the significance of frequencies of leadership-event sequences* we propose a null model of the sequences of leadership events and perform hypothesis testing to compare the support distribution of leadership-event sequences inferred from the given input to our proposed null model.

We use several simulated datasets from the work in Amornbunchornvej and Berger-Wolf (2018a) that cover various types of leadership dynamics for validation, as well as a dataset of trajectories of baboons (Crofoot et al. 2015; Strandburg-Peshkin et al. 2015) to demonstrate the application of our framework. To the best of our knowledge, this is the first work to deal with the topic of complex leadership dynamics and there is no comparable method, and therefore, we indirectly compare our framework to an enhanced FLOCK method (Andersson et al. 2008), used as a baseline for leadership inference only.

1.2 New contributions: followership dynamics

While our previous work in Amornbunchornvej and Berger-Wolf (2018b) addressed several aspects of leadership dynamics, there are still gaps remaining in our understanding of followership dynamics. For instance, assuming individuals i and j are in the same subgroup, how likely is it that they will be in different groups in the future? How many clusters of individuals are there such that members in each cluster stay together with a support at least 0.7? How likely is it for an individual i that its subgroup will be lead by an individual j from the same group?

> Mining Patterns of Followership Dynamics: Given the time series of individual activities, the goal is to mine and model frequent patterns of followership dynamics, including the change of subgroups of followers (unity), or the choice of whom to follow (loyalty).

To address these questions, in this paper, we extend our framework to include several aspects of followership dynamics. In addition to the previous work in Amornbunchornvej and Berger-Wolf (2018b), our framework is capable of:

- *Mining and modeling frequent patterns of faction membership* estimating the frequency of a pair of individuals being in the same faction, as well as discovering faction clusters of individuals.
- *Mining and modeling frequent patterns of leader–follower relationship* estimating the frequency of each individual being a follower of a group lead by a specific individual, as well as discovering dynamics of the changes of a leader of each faction cluster.

We also use simulated datasets from the work in Amornbunchornvej and Berger-Wolf (2018a) that contain complex leader–follower dynamics to evaluate our framework. Our approach is flexible to be generalized beyond the time series of movement data to arbitrary time series where subsets intentionally or spontaneously coordinate.

2 Problem statement

In this paper, we use leadership definitions from the work in Amornbunchornvej and Berger-Wolf (2018a). Given a D-dimensional time series Q, we use $Q(t)$ to refer to an element of the time series Q at time t and, for a given $\Delta \in \mathbb{Z}$, Q_Δ as a time-shifted version of Q where, $Q(t) = Q_\Delta(t + \Delta)$.

Definition 1 [*σ-Following relation* (Amornbunchornvej and Berger-Wolf 2018a)] Let $\mathcal{U}$ be a set of time series, $\mathrm{sim} : \mathcal{U} \times \mathcal{U} \rightarrow [0, 1]$ be a time series similarity function, and $\sigma \in [0, 1]$ be a similarity threshold. For any $P, Q \in \mathcal{U}$, we say that Q *follows* P if Q and P are sufficiently similar within some time shift Δ:

$$\max_{\Delta}(\mathrm{sim}(P, Q_\Delta)) \geq \sigma \text{ and}$$
$$\min(\operatorname*{argmax}_{\Delta} \mathrm{sim}(P, Q_\Delta) \geq 0) \neq \emptyset.$$

Typically, to measure the following relation in Definition 1, the work by Amornbunchornvej and Berger-Wolf (2018a) used the dynamic time warping (DTW) developed by Sakoe and Chiba (1978). DTW is used to measure a distance between two time series. It can measure a distance of multi-dimensional time series (Keogh and Ratanamahatana 2005). Since DTW uses Euclidean distance as a kernel to measure a distance between a pair of elements of two time series, the weighted Euclidean distance can be deployed in the case that we want to give some dimensions higher contribution to a distance measure. Additionally, DTW performs better than several methods (Kjargaard et al. 2013) and robust to the noise (Shokoohi-Yekta et al. 2015) for the task of following relation inference (Amornbunchornvej et al. 2018). The following relation measure using DTW is bounded in [0, 1] interval (Eq. 2). In the case that sim is not bounded, and then, we need a threshold τ to normalize the similarity measure. If $\mathrm{sim} \geq \tau$, then the similarity value is one, otherwise zero.

Definition 2 [*Following network* (Amornbunchornvej and Berger-Wolf 2018a)] Let $\mathcal{U}$ be a set of time series. A digraph $G = (V, E)$ is a *following network* of $\mathcal{U}$ where each node in V corresponds to a time series in $\mathcal{U}$ and $(Q, P) \in E$ if Q follows P.

Definition 3 [*Initiator of faction* (Amornbunchornvej and Berger-Wolf 2018a)] Let $G = (V, E)$ be a following network. $L \in V$ is an *initiator of faction* F_{L} if the out-degree of L is zero and the in-degree of L is greater than zero. The member nodes of F_{L} are any nodes that have a directed path in G to L.

2.1 Leadership dynamics

We create a dynamic following network $\mathcal{G} = \langle G_t \rangle$ by considering each temporal sub-interval t of $\mathcal{U}$ of length ω

(time window parameter) and creating a following network G_t of that interval. We then define the notion of the time series of leaders of a dynamic following network below.

Definition 4 (*Time series of leaders*) Let $\mathcal{U}$ be a set of time series. $\mathcal{L}$ is a *time series of leaders* where $\mathcal{L}(t)$ is a set of faction initiators at time t in G_t.

We can use mFLICA framework (Amornbunchornvej and Berger-Wolf 2018a) to extract a time series of leaders from time series of movement. Next, we define the support of a leader set S. Let T be the length of the time series of leaders and $\mathbb{1}_x$ be an indicator function, which is 1 if the statement x is true, and 0 otherwise.

$$\text{supp}_{\mathcal{L}}(S) = \frac{\sum_{t=1}^{T} \mathbb{1}_{S=\mathcal{L}(t)}}{T}. \tag{1}$$

$\text{supp}_{\mathcal{L}}(S)$ indicates the support of having a particular set of initiators S lead multiple groups concurrently. For example, if $\text{supp}_{\mathcal{L}}(\{L_1, L_2\}) = 0.5$ it means that half the time the leaders are exactly $\{L_1, L_2\}$, leading their factions concurrently.

Definition 5 (*Frequent-leader set*) Let $\mathcal{L}$ be a time series of leaders, S be a set of faction initiators, and $\phi \in [0, 1]$ be a support threshold. S is a frequent-leader set of $\mathcal{L}$ if $\text{supp}_{\mathcal{L}}(S) \geq \phi$.

Definition 6 (*Transition probability of leader sets*) Let $\mathcal{L}$ be a time series of leaders, and S_i, S_j be sets of faction initiators. A transition probability of leader sets λ_{S_i,S_j} is a probability that $\mathcal{L}(t-1) = S_i$ and $\mathcal{L}(t) = S_j$.

Now, we are ready to formally state the problem of Mining Patterns of Leadership Dynamics.

Problem 1: Mining Patterns of Leadership Dynamics

Input: A set $\mathcal{U} = \{U_1, \dots, U_n\}$ of m-dimensional time series and a support threshold ϕ.

Output: A set of frequent-leader sets $\mathcal{S}_{\mathcal{L}}$ and a transition probability set $\mathcal{P} = \{\lambda_{S_i,S_j}\}$ where $S_i, S_j \in \mathcal{S}_{\mathcal{L}}$.

In this paper, we choose to represent a set of frequent-leader sets as a diagram of leadership dynamics below.

Definition 7 (*A diagram of leadership dynamics*) Let $\mathcal{L}$ be a time series of leaders, $\phi \in [0, 1]$ be a support threshold, and $\mathcal{S}_{\mathcal{L}}$ be set of frequent-leader sets. A digraph $\mathcal{T} = (V_{\mathcal{T}}, E_{\mathcal{T}})$ is a diagram of leadership dynamics such that the nodes $V_{\mathcal{T}}$ represent frequent-leader sets $\mathcal{S}_{\mathcal{L}}$ and $(v_i, v_j) \in E_{\mathcal{T}}$ if $\lambda_{S_i,S_j} > 0$.

2.2 Followership dynamics

Given a dynamic following network $\mathcal{G} = \langle G_t \rangle$ of a set of time series $\mathcal{U}$, we define the notion of the *time series of factions* of the dynamic following network below.

Definition 8 (*Time series of factions*) Let $\mathcal{U}$ be a set of time series and $\mathcal{G} = \langle G_t \rangle$ be a dynamic following network of $\mathcal{U}$, where $G_t = (V_t, E_t)$ is a following network at time t. We say that $\mathcal{F}$ is a *time series of factions*, where $\mathcal{F}(t)$ is a set of factions $\{F_{\text{L}}\}$ at time t in $G_t = (V_t, E_t)$ and $F_{\text{L}} \subseteq V_t$ is a set of faction members of initiator L (Definition 3).

The time series of factions contains the information of who belongs to which specific faction over time. Having defined the time series of factions, we can formalize the concept of a pair of individuals who frequently stay together in the same faction.

Definition 9 (*Frequent co-faction pair*) Let $\mathcal{F}$ be a time series of factions, ϕ_{CO} be a threshold, and $V = \{1, \dots, n\}$ be a set of individual indices. We say that individuals i and j in V are a *frequent co-faction pair* if the frequency of i and j being in the same faction of $\mathcal{F}$ is greater than ϕ_{CO}.

A frequent co-faction pair might indicate friendship or other strong affiliation between individuals. At the group level, we define the notion of a frequent co-faction cluster below.

Definition 10 (*Frequent co-faction cluster*) Let $\mathcal{F}$ be a time series of factions, ϕ_{CO} be a threshold, and $V = \{1, \dots, n\}$ be a set of individual indices. We say that a set $C \subseteq V$ is a *frequent co-faction cluster* if every member pair i, j in C is a frequent co-faction pair with respect to $\mathcal{F}$ and ϕ_{CO} and there is no other frequent co-faction cluster $C' \subseteq V$ where $C \subset C'$. In other words, C is a maximal set of frequent co-faction pairs.

A frequent co-faction cluster represents a concept of cohesion. If an entire group has a strong level of cohesion, then there are a few clusters. In contrast, if a group has a weak level of cohesion, then there are multiple clusters. The members within the same cluster might either be loyal to a specific leader, share the same interests, or be a group of friends or other strong affiliates.

In general, given a time series of factions as an input, the problem of finding global faction clusters of individuals is NP-hard, and an approximated solution was given by Tantipathananandh et al. (2007). However, in our setting, a faction or cluster is a group of followers lead by a particular frequent leader. Hence, the problem of finding global faction clusters reduces to the problem of finding followers of each

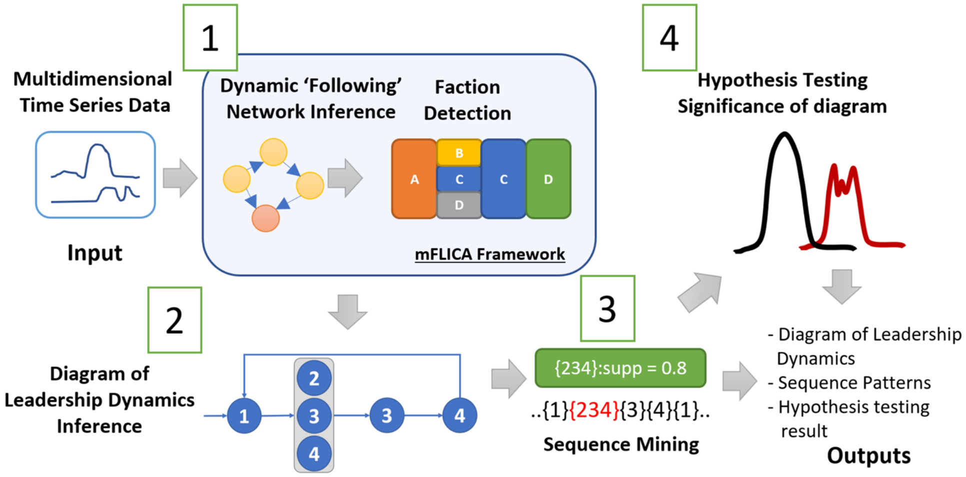
1
Multidimensional Time Series Data
Input
Dynamic 'Following' Network Inference
Faction Detection
A
B
C
D
C
D
mFLICA Framework
4
Hypothesis Testing Significance of diagram
2
Diagram of Leadership Dynamics Inference
1
2
3
4
3
4
3
{234}:supp = 0.8
..{1}{234}{3}{4}{1}..
Sequence Mining
- Diagram of Leadership Dynamics
- Sequence Patterns
- Hypothesis testing result
Outputs

Fig. 1 High-level overview of the proposed framework for inferring leadership dynamics

frequent leader, which can be done in polynomial time in one scan over the time series.

To illustrate the concept of loyalty of a follower toward a specific leader, we formalize the notion of a *frequent leader–follower pair* below.

Definition 11 (*Frequent leader–follower pair*) Let $\mathcal{F}$ be a time series of factions, ϕ_{LF} be a threshold, and $V = \{1, \dots, n\}$ be a set of individual indices. We define the order pair (i, j) in V is the *frequent leader–follower pair* if the frequency of having i within a faction that has j as the initiator with respect to $\mathcal{F}$ is greater than ϕ_{LF}.

A frequent leader–follower pair i, j implies that i is a member of a faction lead by j most of the time. This implies that i might be a loyal follower of j (or have a strong affiliation with a loyal follower of j).

Now, we are ready to formally state the problem of Mining Patterns of Followership Dynamics.

Problem 2: Mining Patterns of Followership Dynamics

Input : A set $\mathcal{U} = \{U_1, \dots, U_n\}$ of m-dimensional time series, threshold ϕ_{CO}, and threshold ϕ_{LF}.

Output: A set of co-faction pairs $\mathcal{S}_{\mathrm{CO}}$, a set of frequent co-faction clusters $\mathcal{S}_{CL}$, and a set of frequent leader–follower pair $\mathcal{S}_{\mathrm{LF}}$.

In this paper, we choose to represent a set of frequent co-faction pairs as a co-faction network, as well as choose to represent a set of frequent leader–follower pairs as a lead–follow network below.

Definition 12 (*A co-faction network*) Let $\mathcal{F}$ be a time series of factions, ϕ_{CO} be a threshold, and $V = \{1, \dots, n\}$ be a set of individual indices. An undirected graph $\mathcal{G}_{\mathrm{CO}} = (V, E_{\mathrm{CO}})$ is a co-faction network such that there is an edge $(v_i, v_j) \in E_{\mathrm{CO}}$ if a pair (v_i, v_j) is a frequent co-faction pair w.r.t. $\mathcal{F}$ and ϕ_{CO}. The weight of the edge (v_i, v_j) is a frequency of having i, j in the same faction.

Definition 13 (*A lead–follow network*) Let $\mathcal{F}$ be a time series of factions, ϕ_{LF} be a threshold, and $V = \{1, \dots, n\}$ be a set of individual indices. A bipartite graph $\mathcal{G}_{\mathrm{LF}} = (V_{\mathrm{F}}, V_{\mathrm{L}}, E_{\mathrm{LF}})$ is a lead–follow network where V_{F} represents a set of follower nodes, and V_{L} represents a set of initiator nodes. For any $v_i \in V_{\mathrm{F}}$ and $v_j \in V_{\mathrm{L}}$, there is a directed edge $(v_i, v_j) \in E_{\mathrm{LF}}$ if an order pair (v_i, v_j) is a frequent leader–follower pair w.r.t. $\mathcal{F}$ and ϕ_{LF}.

3 Methods

3.1 Leadership dynamics

To solve Problem 1, we propose the framework consisting of four parts (Fig. 1). Given a set of time series of movement $\mathcal{U} = \{U_1, \dots, U_n\}$, where $U_i \in \mathcal{U}$ is a two-dimensional time series of length T; first, we infer a dynamic following network and time series of leaders $\mathcal{L}$ using mFLICA framework (Amornbunchornvej and Berger-Wolf 2018a) (Sect. 3.2). Second, we infer a diagram of leadership dynamics $\mathcal{T}$ from $\mathcal{L}$ in Sect. 3.3. Third, we detect the sequence patterns on $\mathcal{L}$ in Sect. 3.4. Finally, we deploy hypothesis tests to evaluate significance of leadership dynamics compared to our proposed null models in Sect. 3.5.

3.2 mFLICA

Given a pair of time series U and Q, mFLICA uses dynamic time warping (DTW) (Sakoe and Chiba 1978) to infer a following relation. Suppose $P_{U,Q}$ is an optimal warping path from DTW dynamic programming matrix, where $(i,j) \in P_{U,Q}$ implies $U(i)$ matched with $Q(j)$ in the matrix. Intuitively, if U is followed by Q with the time delay $\Delta_{i,j}$, then $j - i = \Delta_{i,j}$. Hence, we can compute the *following* relation by the equation below.

$$\mathrm{f}(P_{U,Q}) = \frac{\sum_{(i,j)\in P_{U,Q}} \mathrm{sign}(j-i)}{|P_{U,Q}|}. \tag{2}$$

Suppose we have a similarity threshold σ, then we say there is a following relation if $|\mathrm{f}(P_{U,Q})| \geq \sigma$, where Q follows U if $\mathrm{f}(P_{U,Q}) \geq \sigma$ and U follows Q if $\mathrm{f}(P_{U,Q}) \leq -\sigma$. We set $\sigma = 0.5$ as a default.

Next, given a time window ω and a sliding window parameter $\delta = 0.1\omega$, we have the i^{th} time window interval $w(i) = [(i-1)\times\delta, (i-1)\times\delta+\omega]$. mFLICA creates a following network for each set of time series within interval $w(i)$ of $\mathcal{U}$. An edge of a following network is inferred by Eq. 2 with the weight $|\mathrm{f}(P_{U,Q})|$. Hence, after every interval $w(i)$ has its following network, and we have a dynamic following network $\mathcal{G} = \langle G_t \rangle$ of $\mathcal{U}$.

Lastly, for each time step t, mFLICA uses breadth-first search (BFS) to infer factions and initiators within a following network G_t. The faction initiators are nodes with out-degree zero and in-degree nonzero. By applying BFS to dynamic following network $\mathcal{G}$, we have the time series of leaders $\mathcal{L} = (\mathcal{L}(1), \dots, \mathcal{L}(T))$ as well as the time series of factions $\mathcal{F} = (\mathcal{F}(1), \dots, \mathcal{F}(T))$ as the outputs of this step.

3.3 Inferring transition diagram of leadership dynamics

We use Hidden Markov Model (HMM) (Rabiner 1989) to model a diagram of leadership dynamics $\mathcal{T} = (V_{\mathcal{T}}, E_{\mathcal{T}})$ in Definition 7 and use Baum–Welch algorithm (Jelinek et al. 1975) to infer the maximum likelihood estimates of parameters of HMM from the time series of leader $\mathcal{L}$. In this setting, we have a set of frequent-leader sets $\mathcal{S}_{\mathcal{L}}$ as a set of states in HMM with the support threshold $\phi = 0.01$. In HMM, the stochastic transition matrix A, which has its size $|\mathcal{S}_{\mathcal{L}}| \times |\mathcal{S}_{\mathcal{L}}|$, describes estimated probabilities that a group changes its current set of leaders to another set of leaders (e.g., group merging or splitting). However, since we are interested only in the events of state changes, we ignore the self-transition probability and normalize A to be A^* (Eq. 3), which is the adjacency matrix of $\mathcal{T}$.

Given a time series of leaders $\mathcal{L}$, we can easily infer a set of leader sets $\mathcal{S}_{\mathcal{L}}$. Then, let $\mathcal{S}_{\text{HMM}}$ be a set of states in HMM where $\mathcal{S}_{\text{HMM}}$ and $\mathcal{S}_{\mathcal{L}}$ are in one-to-one correspondence. We represent each state in S_{HMM} as a number in $[1, |\mathcal{S}_{\text{HMM}}|]$, and then, we create $\mathcal{L}_{\text{HMM}}$ by replacing each element in $\mathcal{L}$ with the number of corresponding state in $\mathcal{S}_{\text{HMM}}$. For example, in Fig. 3 (Dynamic Type 1), we have a set of leader sets $\mathcal{S}_{\mathcal{L}}$ and $\mathcal{S}_{\text{HMM}}$, where $\{\text{ID1}\}$, $\{\text{ID2}, \text{ID3}, \text{ID4}\}$, $\{\text{ID3}\}$, and $\{\text{ID4}\}$, in $\mathcal{S}_{\mathcal{L}}$ have corresponding elements in $\mathcal{S}_{\text{HMM}}$ as 1, 2, 3, and 4, respectively.

Initially, we set a stochastic transition matrix $A = \{a_{i,j}\}$ $(i, j$ are the states) and the initial state distribution π_i uniformly. We have the set of observation values $Y = \{1, \dots, |\mathcal{S}_{\text{HMM}}|\}$. In this setting, there is no hidden state since an observation value is an identity of a state. However, in HMM, at any state i, there is a required probability $b_{i,j}$ of observing value j at the state i (typically represented by a matrix $B = \{b_{i,j}\}$). Here, the probability $b_{i,j} = 1$ if $i = j$ and zero otherwise.

We use Baum–Welch algorithm (Jelinek et al. 1975) to infer $A = \{a_{i,j}\}$, and then, we normalize A to create $A^* = \{a^*_{i,j}\}$ by the equation below.

$$a^*_{i,j} = \begin{cases} 0, & i = j \\ \frac{a_{i,j}}{\sum_{k=1,k\neq j}^{|\mathcal{S}_{\text{HMM}}|} a_{i,k}}, & \text{Otherwise.} \end{cases} \tag{3}$$

3.4 Mining sequence patterns of leadership dynamics

After having a diagram of leadership dynamics $\mathcal{T} = (V_{\mathcal{T}}, E_{\mathcal{T}})$, for each pair of nodes $(i,j) \in V_{\mathcal{T}}$, we find a sequence pattern, which is a path $P_{i,j} = (v(1) = i, \dots, v(k) = j)$, where for all $u \in V_{\mathcal{T}}$, $a^*_{v(t-1),v(t)} > a^*_{v(t-1),u}$.

$P_{i,j}$ is an order sequence that the previous state $v(t-1) \in P_{i,j}$ has the highest probability to change to the next consecutive state $v(t) \in P_{i,j}$, given a starting point at i and the final state at j.

Given $A^* = \{a^*_{i,j}\}$ as an adjacency matrix of $\mathcal{T}$, we convert A^* to be $A' = \{a'_{i,j}\}$ where $a'_{i,j} = \frac{1}{a^*_{i,j}}$. Then, we use the standard Dijkstra's algorithm to find the shortest path between every two nodes in A'. Hence, $P_{i,j}$ is the shortest path between i and j in A'.[1] Let v be a number of times that the full sequence of $P_{i,j}$ occurs in $\mathcal{L}$ and N be a number of times that leadership state change happens in $\mathcal{L}$ (e.g., two subgroups merged together, changing the leader), we can find the support of $P_{i,j}$ in the time series of leader $\mathcal{L}$ by the equation below:

$$\mathrm{supp}_{\mathrm{path}}(\mathcal{L}, P_{i,j}) = \frac{v \times (|P_{i,j}| - 1)}{N}. \tag{4}$$

Specifically, v is a number of times that all pairs of nodes $v(t-1), v(t) \in P_{i,j}$ s.t. $v(t-1)$ appear before $v(t)$ in $P_{i,j}$ also appear in $\mathcal{L}$.

3.5 Hypothesis testing

3.5.1 Evaluating the significance of leadership-event order

Given a time series of leaders $\mathcal{L}$ and a diagram of leadership dynamics $\mathcal{T}$ inferred from $\mathcal{L}$, we perform a random

[1] Note, this can be done since the probability condition is independent of each pair and not cumulative over the path

Table 2 Details of nonparametric tests used in this paper

Method	Null hypothesis H_0
Kolmogorov–Smirnov test (Massey 1951)	Two samples are from the same distribution
Wilcoxon rank sum test (Wilcoxon 1945)	
Kruskal–Wallis test (Kruskal and Wallis 1952)	

A significant level has been set at $\alpha = 0.01$ for all experiments

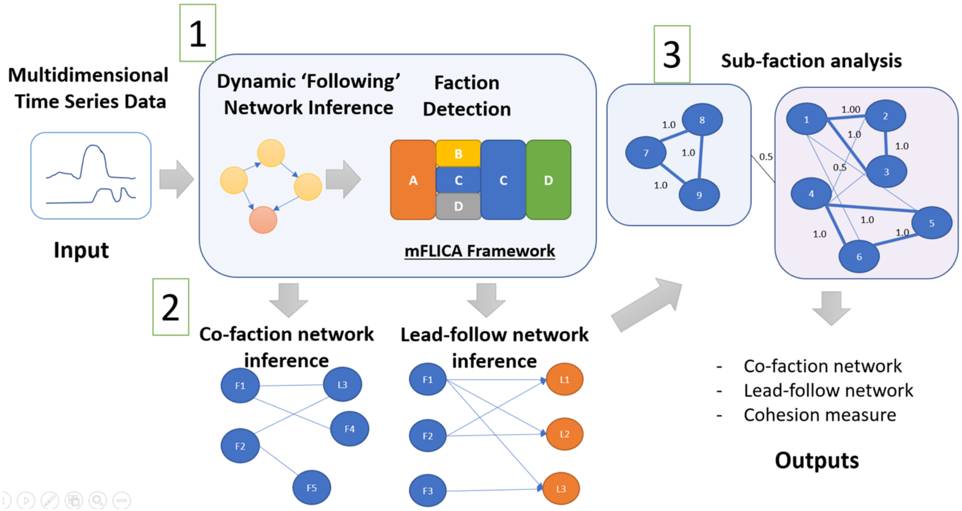

Fig. 2 High-level overview of the proposed framework for inferring followership dynamics

permutation of elements in $\mathcal{L}$ to create $\mathcal{L}_{\text{rand}}$, and then, we infer a diagram of leadership dynamics $\mathcal{T}_{\text{rand}}$ from $\mathcal{L}_{\text{rand}}$ by the method described by the previous section. Afterward, we test the similarity of the edge-weight distributions of $\mathcal{T}$ and $\mathcal{T}_{\text{rand}}$. We deploy three nonparametric methods, shown in Table 2, to perform the tests. If all three methods successfully reject the null hypothesis with the significant threshold $\alpha = 0.01$, then we conclude that the edge-weight distribution of $\mathcal{T}$ is significantly different from $\mathcal{T}_{\text{rand}}$'s although the support value of each node in both graphs is the same.

3.5.2 Evaluating the significance of frequencies of leadership-event sequences

After finding all the sequences for every pair of nodes in Sect. 3.4, we compute the support $\text{supp}_{\text{path}}(\mathcal{L}, P_{i,j})$ of each sequence $P_{i,j}$. This gives the sequence-support distribution of $\mathcal{T}$. Next, we rewire $\mathcal{T}$ to be $\mathcal{T}_{\text{rand}}$ by uniformly and randomly changing the end points of each edge in $\mathcal{T}$, and then, we calculate the sequence-support distribution of $\mathcal{T}_{\text{rand}}$ (Eq. 4). Lastly, we test whether $\mathcal{T}$ and $\mathcal{T}_{\text{rand}}$ sequence-support distributions are different the same way as in the previous section.

We repeat both types of significance tests 100 times and report the percentage of times that the tests successfully reject H_0 for each dataset.

3.6 Followership dynamics

To solve Problem 2, we propose the framework that consists of three parts (Fig. 2).

Let $\mathcal{U} = \{U_1, \dots, U_n\}$ be a set of time series of movement, where $U_i \in \mathcal{U}$ is a time series of length T. In the first step, we infer a dynamic following network as well as a time series of factions $\mathcal{F}$ using mFLICA framework (Amornbunchornvej and Berger-Wolf 2018a) (Sect. 3.2).

Second, we infer a co-faction network $\mathcal{G}_{\text{CO}} = (V, E_{\text{CO}})$ (Sect. 3.6.1) and a lead–follow network $\mathcal{G}_{\text{LF}} = (V, E_{\text{LF}})$ (Sect. 3.6.2) from $\mathcal{F}$. Afterward, we infer a set of frequent co-faction clusters $\{C\}$ (Sect. 3.6.3).

3.6.1 Co-faction network inference

To infer a co-faction network, the first step is to infer a pair of frequent co-faction in Definition 9.

Given a time series of factions $\mathcal{F}$ with the length T and an indicator function $\mathbb{1}_x$ (which is 1 if the statement x is true, and 0 otherwise), we define the support of having individuals i and j in the same faction below.

$$\text{csupp}_{\mathcal{F}}(i,j) = \frac{\sum_{t=1}^{T} \mathbb{1}_{\exists F \in \mathcal{F}(t), \{i,j\} \subseteq F}}{T}. \tag{5}$$

Here, $\mathrm{csupp}_{\mathcal{F}}(i,j)$ indicates the support of having a particular pair of individuals i and j being within the same faction in $\mathcal{F}$.

After we compute the supports csupp for all pairs of individuals, we have a co-faction network $\mathcal{G}_{\mathrm{CO}} = (V, E_{\mathrm{CO}})$. Given a threshold ϕ_{CO}, there is an edge $(v_i, v_j) \in E_{\mathrm{CO}}$ if $\mathrm{csupp}_{\mathcal{F}}(i,j) \geq \phi_{\mathrm{CO}}$. The edge weight of (v_i, v_j) is $\mathrm{csupp}_{\mathcal{F}}(i,j)$.

3.6.2 Lead–follow network inference

To infer a lead–follow network, the first step is to infer a frequent leader–follower pair i, j in Definition 11. Given a time series of factions $\mathcal{F}$ with the length T and an indicator function $\mathbb{1}_x$ (which is 1 if the statement x is true, and 0 otherwise), we can define a support of having individual i in the faction lead by an initiator j below.

$$\mathrm{lfsupp}_{\mathcal{F}}(i,j) = \frac{\sum_{t=1}^{T} \mathbb{1}_{\exists F_j \in \mathcal{F}(t), i \in F_j}}{T}. \tag{6}$$

where F_j is a set of faction members leading by j. Here, $\mathrm{lfsupp}_{\mathcal{F}}(i,j)$ indicates the support of having a particular individual i in the faction leading by an initiator j in $\mathcal{F}$.

After we compute the supports lfsupp for all pairs of individuals, we have a lead–follow network $\mathcal{G}_{\mathrm{LF}} = (V_{\mathrm{F}}, V_{\mathrm{L}}, E_{\mathrm{LF}})$. Given a threshold ϕ_{LF}, for any $i \in V_{\mathrm{F}}$ and $j \in V_{\mathrm{L}}$, there is a directed edge $(v_i, v_j) \in E_{\mathrm{CO}}$ if $\mathrm{lfsupp}_{\mathcal{F}}(i,j) \geq \phi_{\mathrm{LF}}$. The edge weight of (v_i, v_j) is $\mathrm{lfsupp}_{\mathcal{F}}(i,j)$. Higher $\mathrm{lfsupp}_{\mathcal{F}}(i,j)$ implies that there is a higher frequency that i is a member of j's faction. Hence, we can use $\mathrm{lfsupp}_{\mathcal{F}}(i,j)$ as a proxy of loyalty of i to j. Higher $\mathrm{lfsupp}_{\mathcal{F}}(i,j)$ implies that i is more loyal to j.

3.6.3 Clustering and cohesion measure

We use the standard Hierarchical clustering with the shortest distance to link clusters (Sibson 1973) to demonstrate our framework ability. However, any clustering algorithm can be used in our framework to perform the analysis. The Hierarchical clustering algorithm is an agglomerative clustering approach that starts with each individual in a cluster by itself. Then, it keeps merging two closest clusters to be a single new cluster. The algorithm keeps merging on a set of clusters until there is only a single cluster. Given C and C' are clusters and $\mathrm{ADJ}_{\mathrm{CO}}$ is an adjacency matrix of a co-faction network that has its element as $\mathrm{csupp}_{\mathcal{F}}(i,j)$, the distance between two clusters is defined below.

$$\begin{aligned}&\mathrm{dist}_{\mathrm{SingleLink}}(C, C')\\&\quad= \min_{i \in C, j \in C'} (\mathrm{dist}(\mathrm{ADJ}_{\mathrm{CO}}(i,*), \mathrm{ADJ}_{\mathrm{CO}}(j,*)))\end{aligned} \tag{7}$$

where $\mathrm{ADJ}_{\mathrm{CO}}(i,*)$ represents an ith vector row of $\mathrm{ADJ}_{\mathrm{CO}}$ and dist() is a standard euclidean distance. The reason that we compute the distance between the vector of weights of i to all individuals and the vector of weights of j to all individuals in $\mathrm{dist}_{\mathrm{SingleLink}}(C, C')$ is that because two individuals who share the same set of $\mathrm{csupp}_{\mathcal{F}}(i,j)$ are likely members of the same faction. Hence, they should have a small distance.

Next, since there are two types of edges in $\mathrm{ADJ}_{\mathrm{CO}}$: edges that connect members within the same clusters and edges that connect individuals of different clusters. We can use k-means algorithm where $k = 2$ to cluster a list of edge weight of the hierarchical tree into two types: internal edges and external edges. Finally, we link any leaves (individuals) of hierarchical tree that are reachable using internal edges to be a member of the same group to represent a frequent co-faction cluster in Definition 10.

To measure the degree of cohesion of $\mathrm{ADJ}_{\mathrm{CO}}$, we use the standard modularity measure (Q-value) proposed by Newman and Girvan (2004) below.

$$Q(\mathrm{ADJ}, \mathcal{C}) = \sum_{c=1}^{|\mathcal{C}|} (e_{i,i} - a_i^2), \tag{8}$$

where $e_{i,j}$ is a fraction of edges that have one end connected to a node in a cluster i and another end connected to a member of a cluster j, and $a_i = \sum_j e_{i,j}$. The value of Q has a range between -1 and 1. If the value is a large positive, then there are multiple strong clusters; the numbers of edges within groups are greater than the numbers of edges between groups. When there are multiple subgroups that have higher edge weight within the same cluster while edges that connect nodes from different clusters have lower edge weights, then Q is close to one. In contrast, if either there is only one cluster or edge weights of all pairs of nodes are not different from each other, then Q is close to zero. In other words, higher Q implies the higher number of subgroups that have relatively high edge weight between nodes within the same cluster compared to edge weights of nodes from different clusters. Hence, we can use Q as a proxy of cohesion of group. Higher Q implies lower cohesion.

3.7 Time and space complexity

The time complexity of mFLICA is $\mathcal{O}(n^2 \times \omega \times T)$, where n is a number of time series, T is a length of time series, and ω is a time window parameter. The time complexity of Baum–Welch algorithm to infer a diagram of leadership dynamics is $\mathcal{O}(m^2 \times T)$ where m is the number of frequent-leader sets. Typically, $m < n$ since there are fewer frequent-leader sets than individuals. In the followership part, we can scan a time series of factions $\mathcal{F}$ only once to compute everything, which has the time complexity at most $\mathcal{O}(n \times n \times T)$. Hence, our framework's overall time complexity is $\mathcal{O}(n^2 \times \omega \times T)$. For the space complexity, the

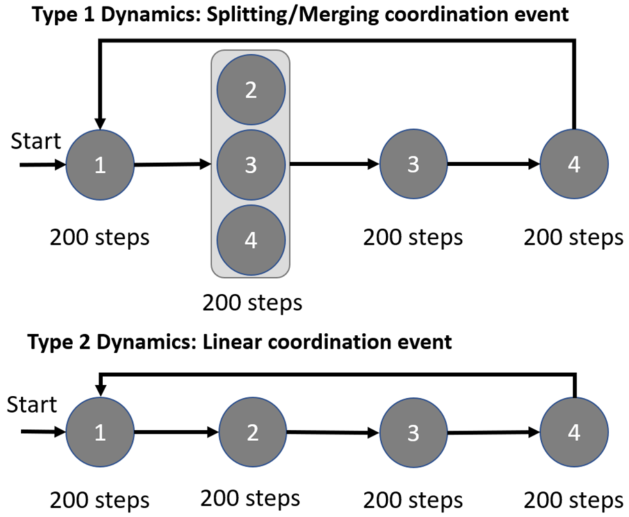
Type 1 Dynamics: Splitting/Merging coordination event
Start
1
2
3
4
3
4
200 steps
200 steps
200 steps
200 steps
Type 2 Dynamics: Linear coordination event
Start
1
2
3
4
200 steps
200 steps
200 steps
200 steps

Fig. 3 Splitting/Merging (above) and linear (below) coordination event. Each node represents the ID of leader of each subgroup at the particular time, and each edge represents the change of group's leaders

most expensive part of our framework is the space for the dynamic following network, which is $\mathcal{O}(n^2 \times T)$.

3.8 Parameters sensitivity

For the time window parameter ω, the work by Amornbunchornvej et al. (2018) reported that the following relation is robust to the noise. However, if we set ω below the maximum time delay between time series, then the result can be severely affected. Hence, a user should try to guess the maximum time delay on his/her dataset before setting ω. Since the core engine of mFLICA is the following relation measure, it is important to set ω properly. The other parameter such as significant level α should be fine-tuned w.r.t. the task.

4 Evaluation datasets

We evaluate our method on synthetic datasets generated using a variety of leadership models with a variety of patterns of leadership dynamics.

4.1 Leadership models

There are three leadership models that we consider in this paper below.

4.1.1 Dictatorship model (DM) (Amornbunchornvej and Berger-Wolf 2018a)

Initially, all individuals stay in the initial area. Then, a single initiator moves toward a target path before others. Afterward, all other individuals follow the initiator with some time delay.

4.1.2 Hierarchical model (HM) (Amornbunchornvej and Berger-Wolf 2018a)

Each individual has been assigned the unique ranking value at the beginning. Lower-rank individuals always follow higher-rank individuals. An initiator who has the highest rank individual (initiator) starts moving first, and then, the second high-rank individual follows the first-rank individual with sometime delay and so on (the $k+1$th rank individual follows the kth-rank individual).

4.1.3 Independent cascade model (IC) (Kempe et al. 2003)

Initially, all individuals are deactivated. At the beginning, each individual has a chance to be active with the probability ρ. After activation, the active individuals move following the initiator except the initiator itself that follows in the target direction. In every time step, active individuals attempt to activate their k-nearest inactive neighbors with the probability of success ρ. Active individuals cannot attempt to activate the same individual twice. In this paper, we determine the parameter space on a combination of: $k \in \{3, 5, 10\}$ and $\rho \in \{0.25, 0.50, 0.75\}$.

4.2 Synthetic trajectory simulation

We use simulated datasets to evaluate the performance of our framework. Each dataset consists of 30 individuals. The trajectory of each individual is two-dimensional time series of length 4000 time steps. Each dataset has been generated from one of the three leadership models described above. There are five coordination events in each dataset. One coordination event lasts for 800 time steps. There are two types of coordination events as follows.

4.2.1 Type 1 dynamics: Splitting/Merging coordination event (Amornbunchornvej and Berger-Wolf 2018a)

In this type of coordination event (Fig. 3), ID1 leads the entire group for 200 time steps. Then, the group splits into three equal size subgroups lead by ID2, ID3, and ID4, for the duration of 200 time steps. Afterward, all subgroups are merged into a single group again lead by ID3 for another 200 time steps. Finally, ID4 leads the entire group for the last 200 time steps.

4.2.2 Type 2 dynamics: Linear coordination event (Amornbunchornvej and Berger-Wolf 2018a)

In this type of coordination event (Fig. 3), ID1 leads first, and then, ID2 leads, ID3 leads after ID2, and ID4 leads after ID3. Each leader leads the group for 200 time steps.

After a coordination event ends, the group stops moving and the next coordination event repeats the pattern. In this paper, we generated 100 datasets for each leadership model and coordination type (e.g., DM with Type 1 dynamics has 100 datasets). One exception: IC has nine cases of different parameters settings, and we have a 100 datasets for each parameter setting and dynamics type. In total, we have 400 datasets for DM and HM but 1800 datasets for IC.

4.3 Baboon dataset

We also deploy our framework on a dataset of GPS trajectories of wild olive baboons (*Papio anubis*) living at Mpala Research Centre, Kenya (Crofoot et al. 2015; Strandburg-Peshkin et al. 2015). The dataset consists of latitude–longitude location time series of 16 baboons recorded for every second for a nine-day period (419,095 time steps). We employ this dataset to demonstrate the potential of our framework to uncover relationships within data to generate scientific hypotheses.

5 Evaluation criteria

5.1 Leadership dynamics

In simulated datasets, we compare the inferred adjacency matrix $A = \{a_{i,j}\}$ of a digraph of leadership dynamics $\mathcal{T} = (V_{\mathcal{T}}, E_{\mathcal{T}})$ against the ground-truth matrix $A^* = \{a^*_{i,j}\}$. For the Splitting/Merging coordination event, the ground-truth set of frequent-leader sets is

$$\mathcal{S}^*_{\mathcal{L}} = \{\{\text{ID1}\}, \{\text{ID2}, \text{ID3}, \text{ID4}\}, \{\text{ID3}\}, \{\text{ID4}\}\}.$$

All elements in A^* are zero except

$$\begin{aligned} a^*_{\{\text{ID1}\},\{\text{ID2,ID3,ID4}\}} &= a^*_{\{\text{ID2,ID3,ID4}\},\{\text{ID3}\}} \\ &= a^*_{\{\text{ID3}\},\{\text{ID4}\}} = a^*_{\{\text{ID4}\},\{\text{ID1}\}} = 1. \end{aligned}$$

For the Linear coordination event,

$$\mathcal{S}^*_{\mathcal{L}} = \{\{\text{ID1}\}, \{\text{ID2}\}, \{\text{ID3}\}, \{\text{ID4}\}\}$$

and all elements in A^* are zero except

$$a^*_{\{\text{ID1}\},\{\text{ID2}\}} = a^*_{\{\text{ID2}\},\{\text{ID3}\}} = a^*_{\{\text{ID3}\},\{\text{ID4}\}} = a^*_{\{\text{ID4}\},\{\text{ID1}\}} = 1.$$

Let $\mathcal{S}_{\mathcal{L}}$ and $\mathcal{S}^*_{\mathcal{L}}$ be the predicted and the ground-truth sets of frequent-leader sets, respectively. The loss function of A and A^* is below:

$$\text{loss}(A, A^*) = \frac{\sum_{i,j \in \mathcal{S}^*_{\mathcal{L}} \cap \mathcal{S}_{\mathcal{L}}} |a_{i,j} - a^*_{i,j}| + \text{FP}(A, A^*) + \text{FN}(A, A^*)}{n_{A^*}} \tag{9}$$

$$\text{FP}(A, A^*) = \sum_{i,j \in \mathcal{S}_{\mathcal{L}} \setminus \mathcal{S}^*_{\mathcal{L}}} |a_{i,j}| \tag{10}$$

$$\text{FN}(A, A^*) = \sum_{i,j \in \mathcal{S}^*_{\mathcal{L}} \setminus \mathcal{S}_{\mathcal{L}}} |a^*_{i,j}| \tag{11}$$

where n_{A^*} is the number of elements within A^*. The first term in Eq. 9 represents the $L1$-norm difference between each element in A and A^* (probabilities) when the predicted states are the same as the ground truth. The second term represents the false positive case when the framework predicts the states that do not exist in the ground truth. The last term represents the false negative case when the framework misses prediction of a state that exists in the ground truth.

5.2 Followership dynamics

5.2.1 Co-faction network

In simulated datasets, we compare an inferred adjacency matrix $A = \{a_{i,j}\}$ of a co-faction network against the ground-truth matrix $A^* = \{a^*_{i,j}\}$. All members from the same cluster are connected with edges that have the weights

$$A^* = \{a^*_{i,j}\} = 1,$$

while two nodes from different clusters have the weight

$$A^* = \{a^*_{i,j}\} = 0.75.$$

For the Splitting/Merging coordination datasets, the ground truth is that there are three clusters:

$$\begin{aligned} C_1 &= \{\text{ID1}, \text{ID3}, \text{ID5}, \dots, \text{ID10}\}, \\ C_2 &= \{\text{ID4}, \text{ID11}, \dots, \text{ID19}\}, \\ C_3 &= \{\text{ID2}, \text{ID20}, \dots, \text{ID30}\}. \end{aligned}$$

For Linear coordination datasets, all individuals are in the single cluster. Given V is a set of nodes of n individuals, we use the absolute loss function to evaluate the difference between predicted $A = \{a_{i,j}\}$ and the ground truth $A^* = \{a^*_{i,j}\}$ below:

$$\text{loss}(A, A^*) = \frac{\sum_{i,j \in V} |a_{i,j} - a_{i,j}|}{\binom{n}{2}}. \tag{12}$$

5.2.2 Lead–follow network

We also compare an inferred adjacency matrix $A = \{a_{i,j}\}$ of a lead–follow network against the ground-truth matrix $A^* = \{a^*_{i,j}\}$ of $\mathcal{G}^*_{\mathrm{LF}} = (V^*_{\mathrm{F}}, V^*_{\mathrm{L}}, E^*_{\mathrm{LF}})$. In both Splitting/Merging and Linear coordination datasets, ID1, ID2, ID3, and ID4 are only initiators. Hence, $V^*_{\mathrm{L}} = \{\mathrm{ID1}, \mathrm{ID2}, \mathrm{ID3}, \mathrm{ID4}\}$ and $V^*_{\mathrm{F}} = \{\mathrm{ID1}, \dots, \mathrm{ID30}\}$.

For Splitting/Merging coordination datasets, given a leader $L = \mathrm{ID1}$, for any $j \in V^*_{\mathrm{F}}, a^*_{\mathrm{L},j} = 0.25$.

- If $L = \mathrm{ID2}$ and $j \in C_3$, then $a^*_{\mathrm{L},j} = 0.25$, while $a^*_{\mathrm{L},j'} = 0$ for $j \notin C_3$.
- If $L = \mathrm{ID3}$ and $j \in C_1$, then $a^*_{\mathrm{L},j} = 0.50$, while $a^*_{\mathrm{L},j'} = 0.25$ for $j \notin C_1$.
- If $L = \mathrm{ID4}$ and $j \in C_2$, then $a^*_{\mathrm{L},j} = 0.50$, while $a^*_{\mathrm{L},j'} = 0.25$ for $j \notin C_2$.

In Linear coordination datasets, for $L \in V^*_{\mathrm{L}}$ and $j \in V^*_{\mathrm{F}}$, $a^*_{\mathrm{L},j} = 0.25$.

Let V_{L} and V^*_{L} be the predicted and the ground-truth sets of initiators of a lead–follow network, respectively, and we compare the inferred A and the ground truth A^* using the loss function below:

$$
\begin{aligned}
&loss_{\mathrm{LF}}(A, A^*) \\
&\quad = \frac{\sum_{i,j \in V^*_{\mathrm{L}} \cap V_{\mathrm{L}}} |a_{i,j} - a^*_{i,j}| + \mathrm{FP}_{\mathrm{LF}}(A, A^*) + \mathrm{FN}_{\mathrm{LF}}(A, A^*)}{n_{A^*}} \\
&\mathrm{FP}_{\mathrm{LF}}(A, A^*) = \sum_{i,j \in V_{\mathrm{L}} \setminus V^*_{\mathrm{L}}} |a_{i,j}| \\
&\mathrm{FN}_{\mathrm{LF}}(A, A^*) = \sum_{i,j \in V^*_{\mathrm{L}} \setminus V_{\mathrm{L}}} |a^*_{i,j}|
\end{aligned} \tag{13}
$$

where n_{A^*} is the number of elements within A^*.

5.2.3 Clustering evaluation

For Splitting/Merging coordination datasets, the ground truth of first cluster is C_1. The second cluster is C_2. The third cluster is C_3. For Linear coordination datasets, all individuals are in the single cluster. We use $F1$ score to measure the difference between inferred and ground-truth clusters. Given C_i is a ground-truth cluster and $\hat{C}_j$ is a predicted cluster that have the most common members with C_i. The true positive is a sum of number of common members between all pair of C_i and $\hat{C}_j$. The false positive is a sum of number of individuals that are in $\hat{C}_j$ but not in C_i, and the false negative is a sum of number of individuals that are in C_i but not in $\hat{C}_j$.

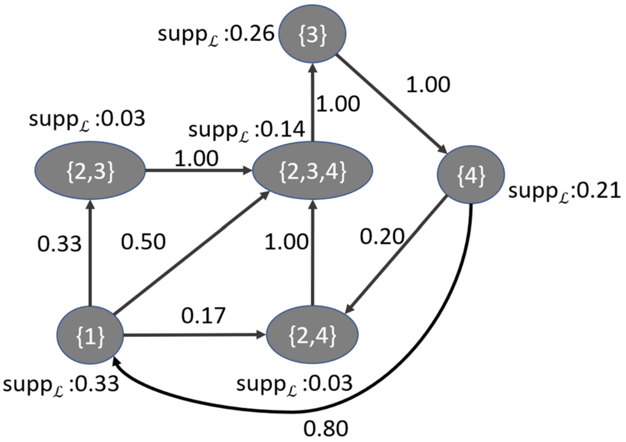
supp$_{\mathcal{L}}$:0.26
{3}
1.00
1.00
supp$_{\mathcal{L}}$:0.03
supp$_{\mathcal{L}}$:0.14
{2,3}
1.00
{2,3,4}
{4}
supp$_{\mathcal{L}}$:0.21
0.33
0.50
1.00
0.20
{1}
0.17
{2,4}
supp$_{\mathcal{L}}$:0.33
supp$_{\mathcal{L}}$:0.03
0.80

Fig. 4 Example of the inferred diagram of leadership dynamics by our framework from a Type-1-HM dataset. Comparing the inferred diagram with the ground truth, only nodes $\{2, 4\}$ and $\{2, 3\}$ are false positive nodes. The support of $\{1\}, \{2, 3, 4\}, \{3\}$ and $\{4\}$ should be 0.25, and our framework can infer the support for each node closely to 0.25

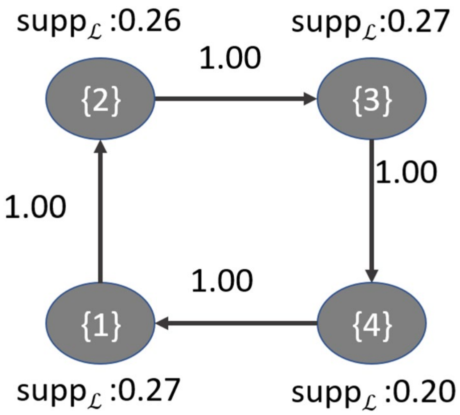
supp$_{\mathcal{L}}$:0.26
supp$_{\mathcal{L}}$:0.27
1.00
{2}
{3}
1.00
1.00
1.00
{1}
{4}
supp$_{\mathcal{L}}$:0.27
supp$_{\mathcal{L}}$:0.20

Fig. 5 Example of the inferred diagram of leadership dynamics by our framework using a Type-2-HM dataset. Comparing the inferred diagram with the ground truth, there are no false positive nodes. The support of $\{1\}, \{2\}, \{3\}$ and $\{4\}$ should be 0.25, and our framework can infer the support for each node closely to 0.25

6 Results

6.1 Leadership dynamics

We set the time window parameter ω using the inference method in Amornbunchornvej and Berger-Wolf (2018a). Figures 4 and 5 show the examples of inferred diagrams of leadership dynamics by our framework from Type-1-HM (Hierarchical model with Splitting/Merging coordination events) and Type-2-HM (Hierarchical model with Linear coordination events) datasets, respectively. In Fig. 4, comparing the inferred diagram with the ground truth, only

Table 3 Example of sequences of leadership dynamics that have the highest support from HM datasets

Datasets	Sequences	$\text{supp}_{\text{path}}(\mathcal{L}, P_{i,j})$
Type-1-HM	{2,3,4},{3},{4},{1}	0.71
Type-2-HM	{1},{2},{3},{4}	0.95

nodes $\{2, 4\}$ and $\{2, 3\}$ are false positive nodes, both with very low support of 0.03.

This implies that despite the complex dynamics of leadership in Type-1 Dynamics case, our framework was still able to retrieve the diagram of leadership dynamics accurately. For the Type-2-HM dataset, which is less complex than Type1-HM case, Fig. 5 shows that there are no false positive nodes in the inferred diagram. Moreover, in both Type-1-HM and Type-2-HM cases, the support of each node should be 0.25, and our framework can infer the support for each node closely to 0.25.

Regarding the mining sequence patterns of leadership dynamics described in Sect. 3.4, Table 3 shows an example of max-support sequences of leadership dynamics that our framework reported from HM datasets. In both dynamics types, the sequences are consistent with the ground truth in Fig. 3.

Next, we compared our framework, which uses the following networks concept (Amornbunchornvej and Berger-Wolf 2018a), to the method based on direction networks proposed in FLOCK method (Andersson et al. 2008) to infer a diagram of leadership dynamics. In direction networks, at any time t, if i is moving toward the same direction as j but j is in front of i, then i follows j. The median of all loss distributions in both Type-1 and Type-2 dynamics datasets is reported in Table 4. The first row of Table 4 shows the distribution of loss values (Eq. 9) in Type-1 dynamics datasets. The direction network approach was reasonably competitive for the Type-1 dynamics. We were able to use the direction networks to infer the states with splits and merges, but the change of leadership was often missed by this underlying method. Not surprisingly, then, the direction network-based method performed significantly worse than the following network-based approach for the Type-2 dynamics. Qualitatively, and as a distribution of the loss values overall, the following networks as the basis for the diagram inference performed better than the direction networks in our framework. In the second row of Table 4, the following networks also performed better than direction networks in Type-2-dynamics datasets.

In Table 5, we reported the hypothesis testing results of the significance of leadership-event order (Sect. 3.5.1). With respect to the type of the leadership model, for the HM, which is a well-structure model, the inferred diagrams are more significantly different from the null-model diagram than for the other leadership models. With respect to the types of the dynamics, in the complex type-1-dynamics datasets our framework inferred diagrams that are more significantly different from the null model. Lastly, the following networks were able to infer diagrams that are more different from the null model than the direction networks.

For hypothesis testing of the significance of frequencies of leadership-event sequences (Sect. 3.5.2), the result is shown in Table 6. Similar to the edge-weight distribution testing, the support distributions of the well-structure model, HM, are significantly different from the support distribution of the null model. The following networks also can be used to infer diagrams that are different from the rewiring diagrams than the direction networks based approach. However, in the simple type-2-dynamics datasets, our framework was able to infer diagrams that are more different from the null model compared to the complex type-1-dynamics case.

For the baboon dataset, we reported the information that we can retrieve from the dataset using our framework as a

Table 4 Median of loss values in the prediction task of diagrams of leadership dynamics

Dyn. type	Type 1			Type 2		
Model	HM	DM	IC	HM	DM	IC
Following network	0.13	0.19	0.24	0	0.03	0.08
Direction network	0.19	0.19	0.25	0.19	0.19	0.25

Table 5 Hypothesis testing results of the significance of leadership-event order in Sect. 3.5.1

Dyn. type	Type 1			Type 2		
Model	HM	DM	IC	HM	DM	IC
Following network	0.99	0.55	0.38	0.86	0.08	0.20
Direction network	0.00	0.00	0.06	0.00	0.00	0.06

We reject H_0 at $\alpha = 0.01$. Each element in the table represents the percentage of the times when the tests successfully reject H_0

Table 6 Hypothesis testing results of the significance of frequencies of leadership-event sequences in Sect. 3.5.2

Dyn. type	Type 1			Type 2		
Model	HM	DM	IC	HM	DM	IC
Following network	0.95	0.35	0.23	0.94	0.84	0.66
Direction network	0.07	0.08	0.20	0.07	0.07	0.20

We reject H_0 at $\alpha = 0.01$. Each element in the table represents the percentage of the case when the test successfully rejects H_0

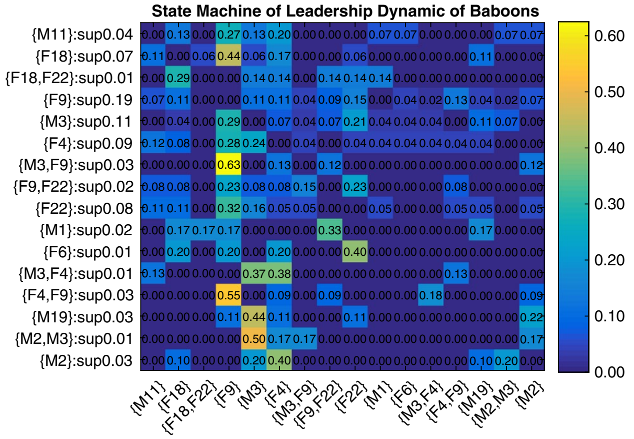

	{M11}	{F18}	{F18,F22}	{F9}	{M3}	{F4}	{M3,F9}	{F9,F22}	{F22}	{M1}	{F6}	{M3,F4}	{F4,F9}	{M19}	{M2,M3}	{M2}
{M11}:sup0.04	0.00	0.13	0.00	0.27	0.13	0.20	0.00	0.00	0.00	0.07	0.07	0.00	0.00	0.00	0.07	0.07
{F18}:sup0.07	0.11	0.00	0.06	0.44	0.06	0.17	0.00	0.00	0.06	0.00	0.00	0.00	0.00	0.11	0.00	0.00
{F18,F22}:sup0.01	0.00	0.29	0.00	0.00	0.14	0.14	0.00	0.14	0.14	0.14	0.00	0.00	0.00	0.00	0.00	0.00
{F9}:sup0.19	0.07	0.11	0.00	0.00	0.11	0.11	0.04	0.09	0.15	0.00	0.04	0.02	0.13	0.04	0.02	0.07
{M3}:sup0.11	0.00	0.04	0.00	0.29	0.00	0.07	0.04	0.07	0.21	0.04	0.04	0.04	0.00	0.11	0.07	0.00
{F4}:sup0.09	0.12	0.08	0.00	0.28	0.24	0.00	0.04	0.00	0.04	0.04	0.04	0.04	0.04	0.04	0.00	0.00
{M3,F9}:sup0.03	0.00	0.00	0.00	0.63	0.00	0.13	0.00	0.12	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.12
{F9,F22}:sup0.02	0.08	0.08	0.00	0.23	0.08	0.08	0.15	0.00	0.23	0.00	0.00	0.00	0.08	0.00	0.00	0.00
{F22}:sup0.08	0.11	0.11	0.00	0.32	0.16	0.05	0.05	0.00	0.00	0.05	0.00	0.00	0.05	0.05	0.00	0.05
{M1}:sup0.02	0.00	0.17	0.17	0.17	0.00	0.00	0.00	0.33	0.00	0.00	0.00	0.00	0.00	0.17	0.00	0.00
{F6}:sup0.01	0.00	0.20	0.00	0.20	0.00	0.20	0.00	0.00	0.40	0.00	0.00	0.00	0.00	0.00	0.00	0.00
{M3,F4}:sup0.01	0.13	0.00	0.00	0.00	0.37	0.38	0.00	0.00	0.00	0.00	0.00	0.00	0.13	0.00	0.00	0.00
{F4,F9}:sup0.03	0.00	0.00	0.00	0.55	0.00	0.09	0.00	0.09	0.00	0.00	0.00	0.18	0.00	0.00	0.00	0.09
{M19}:sup0.03	0.00	0.00	0.00	0.11	0.44	0.11	0.00	0.00	0.11	0.00	0.00	0.00	0.00	0.00	0.00	0.22
{M2,M3}:sup0.01	0.00	0.00	0.00	0.00	0.50	0.17	0.17	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.17
{M2}:sup0.03	0.00	0.10	0.00	0.00	0.20	0.40	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.10	0.20	0.00

Fig. 6 Inferred diagram of leadership dynamics of the baboon dataset from our framework. Each row represents the node of leader sets of the previous state, and each column represents the next state. Each row label consists of baboon gender: 'M' or 'F', a set of frequent-leader IDs, and the support value of frequent-leader set. For example, in row 3 and column 2, the event that two female baboons F18 and F22 are leading their separate subgroups concurrently can happen with the support 0.1 (out of all the coordination times). These two subgroups have a chance to be merged together to a larger group lead by F18 with the probability 0.29.

Table 7 Baboons' sequences of leadership dynamics that have the top-4 highest support

Baboon	Sequences	$\text{supp}_{\text{path}}(\mathcal{L}, P_{i,j})$
Seq. 1	{M11}, {F9}, {M3}	0.0354
Seq. 2	{M18}, {F9}, {M3}	0.0354
Seq. 3	{M18}, {F9}, {M22}	0.0354
Seq. 4	{M4}, {F9}, {M2}	0.0354

case study. Fig. 6 shows the inferred diagram of leadership dynamics from our framework. Each row represents the node of leader sets of the previous state, and each column represents the next state. Each row label consists of baboon gender: 'M' or 'F', a set of frequent-leader IDs, and the support value of frequent-leader set. For example, in row 3 and column 2, the event that two female baboons F18 and F22 are leading their separate subgroups concurrently can happen with the support 0.1 (out of all the coordination times). These two subgroups have a chance to be merged together to a larger group lead by F18 with the probability 0.29. In 4th column ({F9}), we found that no matter what the previous subgroups were, there was a high chance that the next group would be lead solely by the female baboon F9. In 4th row, F9 has the highest support (0.19), which means F9 (who happens to be the dominant female) often leads the troop, with the next highest support of 0.11 for the male baboon M3 (5th column, the alpha male). Lastly, at row 5 and column 4, if M3 and F9 are leading their separate subgroups, then the two groups will be merged to a larger group lead by F9 with probability 0.63.

The hypothesis testing of the edge-weight distribution shows that the baboon's diagram is significantly different from the null model, with 100% of the time the tests successfully rejecting H_0. However, for the hypothesis testing of sequence-support distributions, the baboons' sequences of leadership dynamics are not significantly different from the rewired diagram. Only 5% of the times the tests successfully reject H_0. This indicates that while individual leaders identity is non-random and pairwise leadership transition patterns are significant, there are no leadership sequences that often appear significantly within the baboon dataset. Nevertheless, Table 7 shows baboons' sequences of leadership dynamics that have the top-4 highest support. This result is the evidence that F9 is an important individual who frequently leads the group.

These results show that our framework provides the opportunity for scientists to gain more insight into their datasets in order to generate scientific hypotheses, which might lead to important scientific discoveries (in this case, about the collective behavior and leadership dynamics of social animals).

6.2 Followership dynamics

6.2.1 Co-faction and lead–follow networks

Figure 7 shows the results of ground-truth and predicted adjacency matrices of co-faction network by our framework from Type-1-HM (Hierarchical model with Splitting/Merging coordination events) and Type-2-HM (Hierarchical model with Linear coordination events) datasets. Each

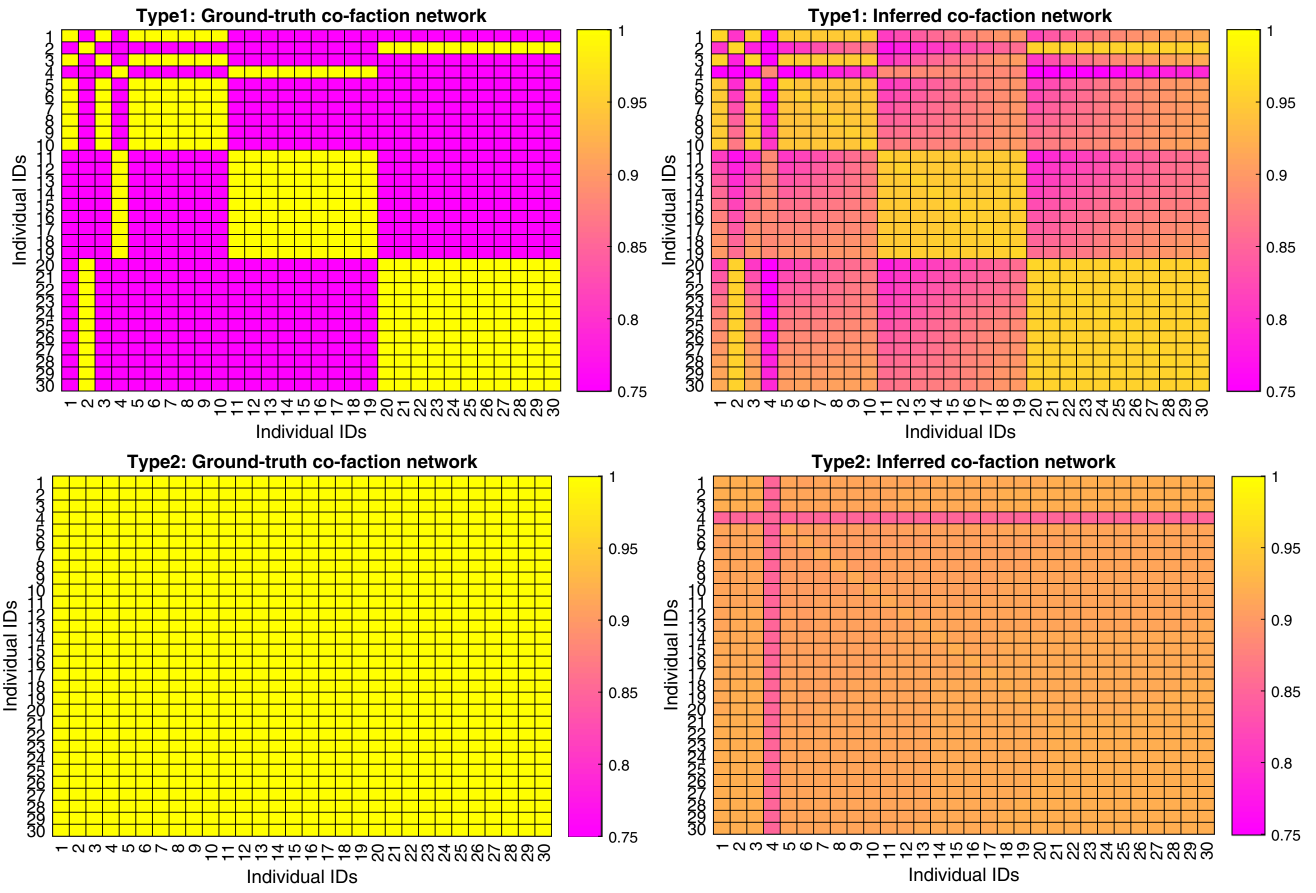

Fig. 7 Adjacency matrices of ground-truth and predicted co-faction networks. (Top-left) Ground-truth matrix of Type-1 dynamics. (Top-right) Predicted matrix of Type-1 dynamics. (Bottom-left) Ground-truth matrix of Type-2 dynamics. (Bottom-right) Predicted matrix of Type-2 dynamics. Each predicted matrix is the result of aggregation of co-faction adjacency matrices from 100 datasets. The lighter color implies a higher value of $\text{csupp}_{\mathcal{F}}(i,j)$

predicted matrix is the result of aggregation of co-faction adjacency matrices from 100 datasets. The result shows that our inferred matrices are mostly similar to the ground-truth matrices with some variations due to noise. ID4 has the highest error in Fig. 7 since it appears during the interval when the group stop moving. Because mFLICA is designed to handle movement initiation analysis, it has a limitation to analyze stopping intervals of movement. Hence, mFLICA cannot capture the behavior of a leader ID4 well.

Figure 8 shows the results of ground-truth and predicted adjacency matrices of lead–follow networks with $\phi_{\text{LF}} = 0.1$. Each predicted matrix is the result of aggregation of lead–follow adjacency matrices from 100 datasets. The result also shows that our inferred matrices are mostly similar to the ground-truth matrices with some variations. The ID4 result has the highest error because of mFLICA limitation that we have just discussed.

We also report the quantitative result of prediction of both co-faction and lead–follow networks using following networks compared with direction networks in Table 8. Overall, our proposed framework using following networks performed better than the direction network framework. For co-faction networks, the loss values are higher than lead–follow network loss values. This implies that finding who are in the same faction frequently is a bit harder than finding who are loyal members of specific leaders.

6.2.2 Clustering results

In the clustering task, given a co-faction network as an input, we compared our proposed framework with a standard community detection algorithm in Newman (2004). The NM community detection method greedily searches for the partition of individuals that maximize the Q-value in Eq. 8. Table 9 shows the result of Q-values of our framework and NM community detection. For type-1 dynamics, we should have a high value of Q-value since there are three strong clusters. Table 9 shows that even

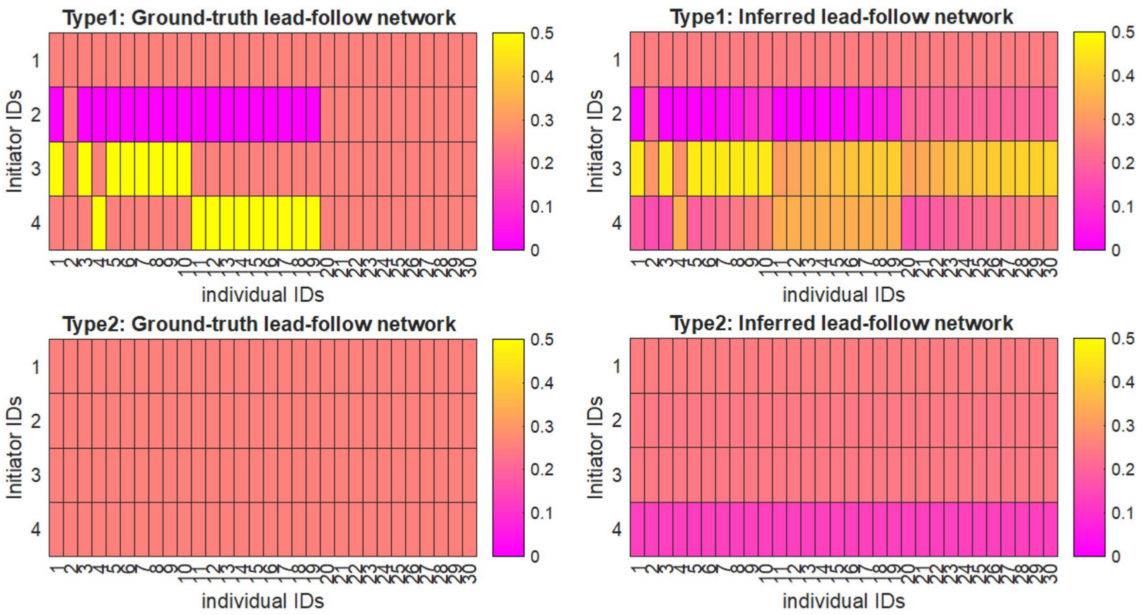

Fig. 8 Adjacency matrices of ground-truth and predicted lead–follow networks. (Top-left) Ground-truth matrix of Type-1 dynamics. (Top-right) Predicted matrix of Type-1 dynamics. (Bottom-left) Ground-truth matrix of Type-2 dynamics. (Bottom-right) Predicted matrix of Type-2 dynamics. Each predicted matrix is the result of aggregation of lead–follow adjacency matrices from 100 datasets. The lighter color implies a higher value of $\mathrm{lfsupp}_{\mathcal{F}}(i, L)$ where i is a column individual (follower) and L is a row individual (initiator)

Table 8 Loss values of co-faction and lead–follow networks inference

	Following Network		Direction Network	
	Type-1	Type-2	Type-1	Type-2
Co-fact loss	0.184 ± 0.013	0.187 ± 0.030	0.398 ± 0.014	0.451 ± 0.011
Lead–foll loss	0.054 ± 0.012	0.026 ± 0.001	0.063 ± 0.007	0.057 ± 0.002

Each element represents a mean of loss value ± two standard deviations from 100 datasets. A lower value implies a better performance of inference

Table 9 Q-value in Eq. 8 of clustering results

	Our framework	NM community detection
Type-1 dynamics	0.6934 ± 0	0.460 ± 0.145
Type-2 dynamics	0.064 ± 0	0.066 ± 0

Each element represents a mean of Q-value value ± two standard deviations from 100 datasets. We expect Type-1 dynamics to have higher Q-value since there are three strong clusters, while Type-2 dynamics should have the Q-value around zero. We report the results from our framework and the candidate approach NM (Newman 2004)

Table 10 F1-score of ground-truth versus inferred clustering results

	Our framework	NM clustering
Type-1 dynamics	0.983 ± 0.003	0.940 ± 0.096
Type-2 dynamics	1 ± 0	0.983 ± 0

Each element represents a mean of F1-score value ± two standard deviations from 100 datasets. A higher F1 score value implies a better performance of clustering inference. We report the results from our framework and the candidate approach NM (Newman 2004)

though NM method tried to find the best clustering partition that maximizes Q-value, our framework found a set of better clusters that has a higher Q-value than NM's clusters. For type-2 dynamics, since there is only one cluster, we expect that the Q-value should be close to zero. Both methods performed well in this case.

We also reported the results of clustering comparison between the ground-truth and inferred clusters. The result in Table 10 shows that our framework performed better than NM in both types of dynamics.

6.3 Baboon followership dynamics

For the baboon dataset, we reported the result of co-faction clustering (Fig. 9) and lead–follow network (Fig. 10)

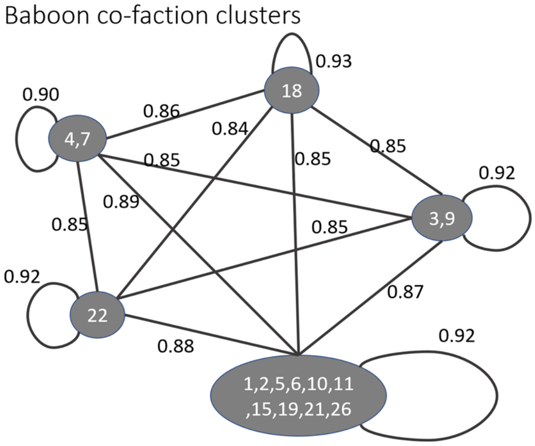

Fig. 9 Co-faction clusters of the baboons dataset inferred by our framework. Each node is a cluster labeled with IDs of cluster members and each edge is a median of $\text{csupp}_{\mathcal{F}}$ of members between clusters

inferred from the trajectories of baboon during pre-coordination intervals of high-coordination events.

Figure 9 shows five major clusters in the dataset. The interesting cluster is the cluster of ID3 and ID9. ID3 is an alpha male while ID9 is an alpha female. Since they are in the same cluster, this implies that they might be a couple.

Figure 9 shows a lead–follow network of the troop. It shows that ID3 and ID9 frequently act as initiators of the group that everyone follows. In both Figs. 9 and 10, we support the hypothesis that ID3 and ID9 might be a center of influence of the group decision-making.

7 Conclusion

In this paper, we proposed a new approach to analyze time series of group movement data. We formalized a new computational problem, Mining Patterns of Leadership Dynamics, and Mining Patterns of Followership Dynamics, as well as proposed a framework as a solution of these problems. Our framework can be used to address several questions regarding leadership and followership dynamics of group movement, such as ‘what is the probability of having two subgroups lead by *i* and *j* being merged together to be a larger group lead by *k* later?’, ‘what is the frequency of having *i* and *k* co-leading their subgroups concurrently?’, and ‘how likely is it that a specific subgroup that *i* is a member will be leading by an individual *j* from the same faction?’. We use the leadership inference framework, mFLICA (Amornbunchornvej and Berger-Wolf 2018a), to infer the time series of leaders and their factions from movement datasets and then propose an approach to mine and model frequent patterns of both leadership and followership dynamics. We evaluate our framework performance by using several simulated datasets, as well as the real-world dataset of baboon movement to demonstrate the applications of our framework. These are novel computational problems and, to the best of our knowledge, there are no existing comparable methods to address them. Thus, we modify and extend an existing leadership inference framework to provide a non-trivial baseline for comparison. Our framework performs better than this baseline in all datasets. Our framework opens the opportunities for scientists to generate testable scientific hypotheses about the dynamics of leadership in movement data.

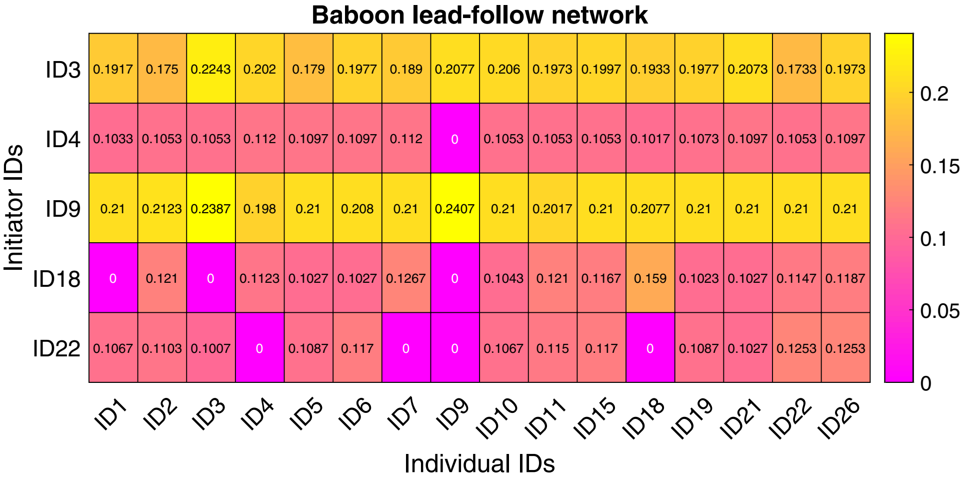

Fig. 10 Lead–follow network of the baboon dataset inferred by our framework

Acknowledgements Funding was provided by National Science Foundation (Grant Nos. III-1514126 and CNS-1248080).

References

Aggarwal CC, Han J (2014) Frequent pattern mining. Springer, Berlin

Agrawal R, Imieliński T, Swami A (1993) Mining association rules between sets of items in large databases. In: Proceedings of the 1993 ACM SIGMOD international conference on management of data, ACM, New York, NY, USA, SIGMOD '93, pp 207–216. https://doi.org/10.1145/170035.170072

Amornbunchornvej C, Berger-Wolf T (2018a) Framework for inferring leadership dynamics of complex movement from time series. In: Proceedings of the 2018 SIAM international conference on data mining, SIAM, pp 549–557. https://doi.org/10.1137/1.9781611975321.62

Amornbunchornvej C, Berger-Wolf TY (2018b) Mining and modeling complex leadership dynamics of movement data. In: ASONAM'18, IEEE, pp 447–454. https://doi.org/10.1109/ASONAM.2018.8508310

Amornbunchornvej C, Brugere I, Strandburg-Peshkin A, Farine DR, Crofoot MC, Berger-Wolf TY (2018) Coordination event detection and initiator identification in time series data. ACM Trans Knowl Discov Data 12(5):53:1–53:33. https://doi.org/10.1145/3201406

Andersson M, Gudmundsson J, Laube P, Wolle T (2008) Reporting leaders and followers among trajectories of moving point objects. GeoInformatica 12(4):497–528

Couzin ID, Krause J, Franks NR, Levin SA (2005) Effective leadership and decision-making in animal groups on the move. Nature 433(7025):513–516

Crofoot MC, Kays RW, Wikelski M (2015) Data from: shared decision-making drives collective movement in wild baboons. Movebank Data Repository. https://doi.org/10.5441/001/1.kn0816jn

Glowacki L, von Rueden C (2015) Leadership solves collective action problems in small-scale societies. Philos Trans R Soc B 370(1683):20150010

Han J, Cheng H, Xin D, Yan X (2007) Frequent pattern mining: current status and future directions. Data Min Knowl Discov 15(1):55–86. https://doi.org/10.1007/s10618-006-0059-1

Hogg MA (2001) A social identity theory of leadership. Personal Soc Psychol Rev 5(3):184–200

Jelinek F, Bahl L, Mercer R (1975) Design of a linguistic statistical decoder for the recognition of continuous speech. IEEE Trans Inf Theory 21(3):250–256

Kempe D, Kleinberg J, Tardos É (2003) Maximizing the spread of influence through a social network. In: Proceedings of the ninth ACM SIGKDD, ACM, pp 137–146

Keogh E, Ratanamahatana CA (2005) Exact indexing of dynamic time warping. Knowl Inf Syst 7(3):358–386

Kjargaard MB, Blunck H, Wustenberg M, Gronbask K, Wirz M, Roggen D, Troster G (2013) Time-lag method for detecting following and leadership behavior of pedestrians from mobile sensing data. In: Proceedings of the IEEE PerCom, IEEE, pp 56–64

Krause J, Hoare D, Krause S, Hemelrijk C, Rubenstein D (2000) Leadership in fish shoals. Fish Fish 1(1):82–89

Kruskal WH, Wallis WA (1952) Use of ranks in one-criterion variance analysis. J Am Stat Assoc 47(260):583–621. https://doi.org/10.1080/01621459.1952.10483441

Lee JG, Han J, Whang KY (2007) Trajectory clustering: a partition-and-group framework. In: Proceedings of the 2007 ACM SIGMOD international conference on management of data, ACM, pp 593–604

Li Z, Lee JG, Li X, Han J (2010) Incremental clustering for trajectories. In: Kitagawa H, Ishikawa Y, Li Q, Watanabe C (eds) Database systems for advanced applications. Springer, Berlin, pp 32–46

Massey FJ Jr (1951) The Kolmogorov–Smirnov test for goodness of fit. J Am Stat Assoc 46(253):68–78. https://doi.org/10.1080/01621459.1951.10500769

Newman MEJ (2004) Fast algorithm for detecting community structure in networks. Phys Rev E 69:066133. https://doi.org/10.1103/PhysRevE.69.066133

Newman MEJ, Girvan M (2004) Finding and evaluating community structure in networks. Phys Rev E 69:026113. https://doi.org/10.1103/PhysRevE.69.026113

Rabiner LR (1989) A tutorial on hidden markov models and selected applications in speech recognition. Proc IEEE 77(2):257–286. https://doi.org/10.1109/5.18626

Sakoe H, Chiba S (1978) Dynamic programming algorithm optimization for spoken word recognition. IEEE Trans Acoust Speech Signal Process 26(1):43–49

Shokoohi-Yekta M, Wang J, Keogh E (2015) On the non-trivial generalization of dynamic time warping to the multi-dimensional case. In: SDM'15, pp 289–297

Sibson R (1973) Slink: an optimally efficient algorithm for the single-link cluster method. Comput J 16(1):30–34. https://doi.org/10.1093/comjnl/16.1.30

Smith JE, Estrada JR, Richards HR, Dawes SE, Mitsos K, Holekamp KE (2015) Collective movements, leadership and consensus costs at reunions in spotted hyaenas. Anim Behav 105:187–200

Spiliopoulou M, Ntoutsi I, Theodoridis Y, Schult R (2006) Monic: modeling and monitoring cluster transitions. In: Proceedings of the 12th ACM SIGKDD international conference on knowledge discovery and data mining, ACM, New York, NY, USA, KDD '06, pp 706–711, https://doi.org/10.1145/1150402.1150491

Strandburg-Peshkin A, Farine DR, Couzin ID, Crofoot MC (2015) Shared decision-making drives collective movement in wild baboons. Science 348(6241):1358–1361. https://doi.org/10.1126/science.aaa5099

Stueckle S, Zinner D (2008) To follow or not to follow: decision making and leadership during the morning departure in chacma baboons. Anim Behav 75(6):1995–2004

Tantipathananandh C, Berger-Wolf T, Kempe D (2007) A framework for community identification in dynamic social networks. In: Proceedings of the 13th ACM SIGKDD international conference on knowledge discovery and data mining, ACM, New York, NY, USA, KDD '07, pp 717–726. https://doi.org/10.1145/1281192.1281269

Wilcoxon F (1945) Individual comparisons by ranking methods. Biometr Bull 1(6):80–83