

Sequential Dynamic Leadership Inference Using Bayesian Monte Carlo Methods

Qing Li, *Student Member, IEEE*, Bashar I. Ahmad, *Member, IEEE*, Simon J. Godsill, *Senior Member, IEEE*

Abstract—Hierarchy and leadership interactions commonly occur in animal groups, crowds of people and in vehicle motions. Such interactions are often affected by one or more individuals who possess key domain information (e.g. final destination, environmental constraints and best routes) or pertinent traits (e.g. better navigation, sensing and decision making capabilities) compared with the rest of the group. This paper presents a framework for the automatic identification of group structure and leadership from noisy sensory observations of tracked groups. Accordingly, a new leader-follower model is developed which assumes the dynamics of the group to be a multivariate Ornstein–Uhlenbeck process with the designated leader(s) drifting to the destination and followers reverting to the leaders’ state. Sequential Monte Carlo (SMC) approaches, and specifically the sequential Markov chain Monte Carlo (SMCMC) approach, are adopted to infer, probabilistically, the evolving leadership structure. A Rao-Blackwellisation scheme is employed such that the kinematic state of the objects in the group is inferred in closed form by Kalman filtering. Experiments show that the proposed techniques can successfully determine the leadership structures in challenging scenarios with a corresponding enhancement in tracking accuracy through direct consideration of the leadership interactions of the group.

Index Terms—leader-follower model, sequential Markov chain Monte Carlo, group tracking, network tracking

I. INTRODUCTION

IN group tracking scenarios such as for animals, drone swarms and crowds of people, members of a group can be constantly interacting with one another. This collective behaviour was first investigated by Reynolds in [25], where a distributed behavioral model was proposed to describe the Boids flock motion under the rules of centering, collision avoidance, and velocity matching. Besides these general rules, there exist many unique interaction patterns in groups (see e.g., [14, 34]), and a very common one is leadership or hierarchical structure, as is widely observed in nature. A well known example is bird flocking, where for instance geese fly in a V-formation and take turns to lead the group during migration from north to south [6, 24]. A swarm of drones can also exhibit a leader-follower formation where a drone piloted by a human or following a predefined trajectory leads the other swarm members when carrying out missions such as monitoring or delivery [28, 31]. Other typical scenarios include a flight of aircraft and convoy of ground vehicles [23].

The benefits of studying this leader-follower formation in group tracking domain are twofold. First, learning this group

structure information can enhance the tracking performance by incorporating interactions into the design of dynamic models. Second, the interaction pattern itself facilitates higher level understanding of the observed scene, which can be valuable in applications such as anomaly detection, intent prediction for surveillance and crowd management during emergencies. This paper addresses the problem of sequentially inferring the dynamic leadership structure of tracked groups of objects over time.

A. Related Work

The leader-follower strategy has originally been described by Stackelberg in economics for static two-person games with one player being the leader and the other being follower [33]. Later, it has been widely adopted in formation control of a leader-follower robotic/drone system, where one robot/drone is chosen as the leader which controls the movement of the whole group, and the others are the followers with desired relative separations and bearings [5, 31]. The leader-follower model has also found its application and been developed further in the tracking domain [3, 10, 19, 23]. Compared with such prior work, here a new leader-follower model is proposed in which the leadership structure can dynamically evolve over-time, and an online inference routine is sought in lieu of batch inference methods when all relevant sensory data is available.

In the leader-follower model, the followers typically revert to the kinematic state of a *virtual* leader which is an explicit additional group state normally modelled as the average position and velocity of all objects [18, 23, 26]. Apart from virtual leader, the defined leaders could also be the physical entities in the group, for example pigeon(s) in a flock of pigeons or fish in a fish school (i.e. group members). Such leader-follower model has been proposed in [16] for tracking a group of manoeuvring objects, which focuses exclusively on enhancing the group tracking performance by adopting the premise that the identity of the member(s) leading the group is available. By contrast, in this paper, the group leaders are *a priori* unknown physical entities, and the main objective is to sequentially reveal the leadership and dominance hierarchy in the tracked group from the available (noisy) sensory observations with suitable sequential algorithms.

Several methods have been introduced to discover the leadership pattern of the group, e.g. [3, 4, 20]. For instance, in [20], a sparse structure inference technique has been proposed to find the leadership and interaction pattern among group members. A causality reasoning framework has been used to rank objects with respect to their dominance effect [4].

Q. Li, B. I. Ahmad, and S. J. Godsill are with the Engineering Department, University of Cambridge, Cambridge CB2 1PZ, U.K. e-mail: ql289@cam.ac.uk; bia23@cam.ac.uk; sjg30@cam.ac.uk

B. I. Ahmad is with Aveillant, Thales Land and Air Systems, Evolution Business Park, Cambridge CB24 9NG, e-mail: bashar.ahmad@aveillant.com

These two approaches and others, e.g. [3, 8], assume that the interaction pattern within groups is fixed over time and/or apply batch inference when all related sensor measurements of the group are available. Whereas, here the leadership and dominance hierarchy can dynamically change over time and an online inference framework is sought, i.e. leadership is sequentially estimated at the arrival of new observation(s), unlike in [3, 4, 8, 20].

The Bayesian leadership inference framework proposed in this paper builds on and generalises the preliminary work of [17], in which a rotated leadership model was introduced to determine a dynamically changing leadership structure and implemented with sequential Monte Carlo methods. In the current work, both the leader and the number of leaders may vary over time, contrary to the formulation in [17], which stipulates that the group is led by only one individual. Such leadership and dominance patterns are prevalent in animal groups, where the group is led by one or more (usually a small proportion) of informed and/or competent individuals who possess information about the group objectives such as the position of the food source or the migration path. An example explored in this paper is the behaviour of a golden shiner fish school, where in the group a selected few members are trained to know the food source and they are typically followed by the other (non-trained) fish to the food location (i.e. the group final destination) as reported in [29].

Compared to the rotated leadership dynamic model in [17], here we facilitate the task of tracking and leadership inference by exploiting the existence of available prior information that influences the leaders' long-term behaviour, for instance a known endpoint in the case of fish seeking food or homing pigeon flocks. For these tasks, we propose a destination-reverting model [1] to model the leaders' dynamics, such that the motion of the whole group is treated as an Ornstein–Uhlenbeck process with leader(s) reverting to the desired destination and the followers drifting towards the leaders. Such intent-driven leadership behaviour was not addressed in [17] and we employ real data from two animal group tracking scenarios in order to demonstrate that such knowledge can be effectively captured within the proposed framework.

B. Contributions

A principal contribution in this paper is the development of a novel and more flexible leader-follower model based on a rotated leadership formulation. This enables a dynamically changing group leadership structure where one or more group members can be leaders. Within the proposed Bayesian framework, intent or destination-driven behaviour can be also incorporated for cases when the destination is known as *a priori*, and this is modelled by an Ornstein–Uhlenbeck process, for example with leader(s) reverting to the intended endpoint. We note that a special case of our model sets the destination-driving force to zero (for scenarios where destination is unknown or unimportant), and in that case the model for leader(s) reduces to a dynamic leadership version of [17]. An appropriate online inference approach, namely the sequential Markov chain Monte Carlo approach, is then

utilised to infer the leadership state sequentially over time from the available noisy sensory observations. Specifically, a Rao–Blackwellisation scheme is employed such that the kinematic state of the objects can be marginalised out and its posterior distribution inferred using Kalman filtering. We propose several different versions of the SMCMC algorithm here and compare with other more standard sequential Monte Carlo techniques. Finally, real data from two animal group tracking scenarios, as well as synthetic data, are used to illustrate the efficacy of the introduced leadership inference framework. The results from the analysis of real data demonstrate that understanding or validating hypotheses about animal behaviour (e.g. from a biological perspective) can be effectively achieved by using an appropriate Bayesian inference approach, thereby potentially reducing the manual analysis of the recorded data.

C. Paper Outline

The remainder of this paper is structured as follows. Section II presents the leader-follower model for describing group motions and interactions. Section III presents the Rao–Blackwellisation scheme and the inference algorithms are detailed in Section IV. In Section V, the performance of the proposed model is evaluated using both synthetic and real pigeon flock and fish tracking data. Conclusions are given in Section VI.

NOMENCLATURE

α	scale parameter that qualifies the leaders' impact on the position of the follower
β	scale parameter that qualifies the leaders' impact on the velocity of the follower
η	a parameter that quantifies the influence of the endpoint on the leader
γ	the drag coefficient preventing the velocities drifting to large values over time
ρ	the Metropolis-Hastings acceptance ratio
τ	the time interval between successive observations
A_{L_t}	the transition matrix given the current leadership structure L_t at time t
D	the position of the destination in any corresponding dimensionality
d	dimension of the space (e.g., motion in 2-D has $d = 2$)
$F_{D,j}(t)$	an attraction force exerted by the endpoint on the leader j at time t
$F_{L_t,i}(t)$	an attraction force exerted by all the leaders in the leadership structure L_t on the i^{th} follower at time t
K	the total number of all possible leadership structures
L_t	the group leadership structure state at time t
M	the total number of iteration
N	the total number of objects in a group
n	the observation time step at t_n
N_p	number of particles used in the Monte Carlo methods
S	the set of all possible leadership structures
S_k	the k -th leadership structure
t_n	time instant at time step n
$w_n^{(i)}$	the weight of particle i at time step n in the MPF algorithm
Z_n	observations received at time step n

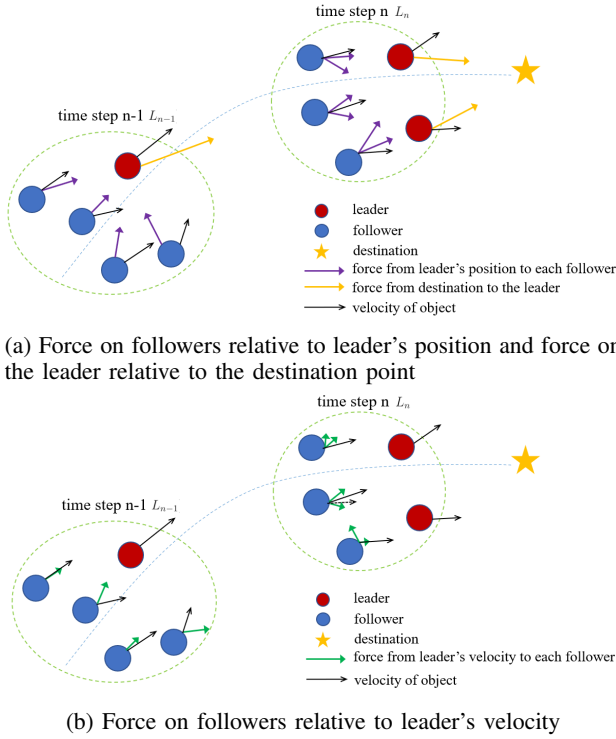


Fig. 1: The dynamic leadership model in the group

II. PROPOSED MODEL

The aim of modelling is to describe the motion of a group led by one or a few group members who possess information about the group intent in terms of destination. Within the introduced framework the term 'destination' has some generality as it can include terminal velocity as well as the more common concept of physical location. In our subsequent formulations, the latter concept is illustrated. It is emphasised that the presented modeling here also applies when the leader(s) behaviour is not driven by an intended endpoint by setting the force destination exerted on leader(s) to zero.

Assume that there are N objects in a group. Let S be the set of all possible group/leadership structures such that $S = \{S_k\}_{k=1}^K$ and $K = |S|$ is their total number. Each S_k specifies the indices of the individuals having a leadership role in the k -th leadership structure. The value of K depends on the prior information about the group structure we process. If no prior information is obtained, then $K = 2^N - 1$ where S lists all possible combinations of each individual being either the leader or the follower (excluding the case with no leadership pattern in the group). For example, if we have a group of 2 objects, then S has a total of $K = 3$ possible leadership structures, with $S_1 = \{1\}$, $S_2 = \{2\}$, $S_3 = \{1, 2\}$. Otherwise the value of K could be much smaller if an informative prior is sought, e.g. an exactly-one-leader condition.

We propose to use two kinds of force to represent group interactions. One set of interactions is between a destination and each leader, such that an attraction force is exerted by this endpoint on the leader with its strength proportional to the distance from the leader's position to the destination. Let $L_t \in S$ be a random variable representing the group leadership

structure state at time t . Assume the motion of the objects in each dimension are mutually independent. For each leader $j \in L_t$ in the d -th dimension, the aforementioned attraction force $F_{D,j}^d(t)$ is expressed as

$$F_{D,j}^d(t) = \eta(D^d - x_j^d(t)), \quad (1)$$

where D^d denotes the position of the destination in any corresponding dimensionality, $x_j^d(t)$ is the position of the leader j , and η is a parameter which quantifies the influence of the intended endpoint on the leader.

The second level of interaction is between leaders and followers. For each follower, a force $F_{L_t,i}^d(t)$ is defined to quantify the influence of all the leaders in the leadership structure L_t on the i th follower at the current time t as per

$$F_{L_t,i}^d(t) = \sum_{j \in L_t} \alpha(x_j^d(t) - x_i^d(t)) + \sum_{j \in L_t} \beta(\dot{x}_j^d(t) - \dot{x}_i^d(t)), \quad (2)$$

where $x_i^d(t)$ and $\dot{x}_i^d(t)$ are the position and velocity of follower i , respectively, and $i \in \{1, \dots, N\} - L_t$, i.e. object i is not a leader. Parameters α and β set the strength of the leader impact on the position and velocity of the i th follower. Here, we assume that the values of α and β are known scalars and the influence of each leader on the followers is identical, for reasons of simplicity. These restrictions are easily relaxed within our framework and parameter learning for these models is outside the scope of this paper.

Consequently, the motion of the group can be described by a set of coupled Ornstein-Uhlenbeck processes governed by two kinds of force, one that tends to drive each leader towards the destination while the other steers the followers towards the leaders' position and velocity. This modelling approach can be readily adapted to scenarios without the destination prior information by simply setting the force $F_{D,j}^d(t)$ to zero. Overall such a dynamic system can be expressed using a continuous-time stochastic differential equation (SDE) as follows

$$d\dot{x}_i^d(t) = (F_{L_t,i}^d(t) - \gamma\dot{x}_i^d(t)) dt + dB_i^d(t), \quad i \in \{1, \dots, N\} - L_t \quad (3)$$

$$d\dot{x}_j^d(t) = (F_{D,j}^d(t) - \gamma\dot{x}_j^d(t)) dt + dB_j^d(t), \quad j \in L_t, \quad (4)$$

where γ is the drag coefficients, preventing the velocities drifting to large values over time. Whereas, $B_i^d(t)$ and $B_j^d(t)$ are independent Brownian motion for each follower and leader.

For each dimension, define the joint state of N objects as $X^d(t) = [x_{1:N}^d(t), \dot{x}_{1:N}^d(t)]^T$. Conditioned on the leadership state L_t , the dynamic model of the joint state $X^d(t)$ is linear and Gaussian. It is modelled by the SDE

$$dX^d(t) = A_{L_t}X^d(t)dt + \eta W dt + C dB^d(t), \quad (5)$$

where $B^d(t) \in \mathbb{R}^N$ is a Brownian motion with the diagonal covariance matrix $Q_w^d = \text{diag}(\sigma_1^2, \sigma_2^2, \dots, \sigma_N^2)$ and matrix $C \in \mathbb{R}^{2N \times N}$ has the form

$$C = \begin{bmatrix} \mathbf{0}_{N \times N} \\ \mathbf{I}_{N \times N} \end{bmatrix}. \quad (6)$$

The transition matrix A_{L_t} and W are given by

$$A_{L_t} = \begin{bmatrix} \mathbf{0}_{N \times N} & \mathbf{I}_{N \times N} \\ \alpha_{N \times N} & \beta_{N \times N} - \gamma_{N \times N} \end{bmatrix}, W = \begin{bmatrix} \mathbf{0}_{N \times 1} \\ D_{N \times 1} \end{bmatrix}, \quad (7)$$

where $\gamma = \text{diag}(\gamma, \dots, \gamma)$. The values of $\alpha_{N \times N}$, $\beta_{N \times N}$ in matrix A_{L_t} , and $D_{N \times 1}$ in vector W depends on the leadership state L_t . Next, those matrices/vectors are defined explicitly.

In matrix $\alpha_{N \times N}$, for each leader j : 1) the corresponding value of $\alpha_{j,j}$ is set as $-\eta$ and 2) other entries in the j^{th} row are all zeros. For each follower i (of a leader j): 3) all entries $\{\alpha_{i,j}\}_{j \in L_t}$ in row i are set as α , 4) $\alpha_{i,i}$ is equal to $-\sum_{j \in L_t} \alpha$, and 5) the other elements in the i^{th} row of α are zeros. In matrix $\beta_{N \times N}$, for each leader j , the corresponding j row is set to the zeros vector $\mathbf{0}_{1 \times N}$. For the follower i , the i^{th} row in $\beta_{N \times N}$ is specified similar to that of $\alpha_{N \times N}$; however α is replaced by β . In vector $D_{N \times 1}$, for each leader j , the corresponding j row equals one, with the other rows for followers being zeros. We recall that α , β and γ are assumed to be known constants.

To clarify the matrices definitions in (7), we present an example for a group of four object with the (inferred) leaders being $\{1, 2\}$. In this case, $\alpha_{4 \times 4}$ is given by

$$\alpha = \begin{bmatrix} -\eta & 0 & 0 & 0 \\ 0 & -\eta & 0 & 0 \\ \alpha & \alpha & -2\alpha & 0 \\ \alpha & \alpha & 0 & -2\alpha \end{bmatrix}, \quad (8)$$

matrix $\beta_{N \times N}$ has a similar form as per

$$\beta = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ \beta & \beta & -2\beta & 0 \\ \beta & \beta & 0 & -2\beta \end{bmatrix}, \quad (9)$$

and vector $D_{N \times 1}$ in vector W in (7) given by

$$D_{4 \times 1} = \begin{bmatrix} D^d \\ D^d \\ 0 \\ 0 \end{bmatrix}. \quad (10)$$

Therefore, the state-space model can be obtained by integrating the SDE in (5)

$$X^d(t + \tau) = F_{L_t} X^d(t) + \mu_D + w_t^d, \quad (11)$$

with the transition matrix F_{L_t} and vector μ_D being

$$F_{L_t} = e^{\tau A_{L_t}}, \quad (12)$$

$$\mu_D = \eta A_{L_t}^{-1} (e^{\tau A_{L_t}} - \mathbf{I}) W, \quad (13)$$

where τ is the time interval between successive observations, which could be time-varying. Here, equally spaced observations are considered as a simplification, which can be easily relaxed for irregular data arrival times.

The dynamical noise in the state space model in (11) is Gaussian, $w_t^d \sim \mathcal{N}(0, P_\tau)$, with covariance P_τ ,

$$P_\tau = e^{\tau A_{L_t}} P_0 e^{\tau A_{L_t}^T} + \int_0^\tau e^{t A_{L_t}} C Q_w^d C^T e^{t A_{L_t}^T} dt. \quad (14)$$

This integral can be evaluated using the matrix fraction decomposition [30]. We especially let $P_0 = C_0 E_0$, $C_0 = P_0 = 0$, $E_0 = \mathbf{I}$ and subsequently we have

$$\begin{pmatrix} C_\tau \\ E_\tau \end{pmatrix} = \exp \left\{ \begin{pmatrix} A_{L_t} & L_t Q_w^d L_t^T \\ 0 & -A_{L_t}^T \end{pmatrix} \tau \right\} \begin{pmatrix} C_0 \\ E_0 \end{pmatrix}, \quad (15)$$

from which the covariance can be calculated by $P_\tau = C_\tau E_\tau$.

Additionally, a linear Gaussian observation model is employed in this paper. Since the focus here is on modelling and then sequentially inferring the group interaction/leadership structure, we assume for simplicity that data association is known *a priori*. This corresponds to the real data examples in Section V-B, where the measurements association has already been manually extracted from video data (e.g., the fish school scenario) or each target carries a sensor (e.g., each pigeon in a flock is equipped with a GPS receiver). For strategies of addressing the data association problem, see for instance [22, 32]. At observation time instant t_n we have

$$Z_n^d = H X_n^d + v_n^d, \quad (16)$$

where X_n^d is the latent state at the measurement time t_n ; n denotes the time step. The i.i.d observation noise obeys $v_n^d \sim \mathcal{N}(0, Q_v)$, with the mapping matrix H defined by

$$H = \begin{bmatrix} \mathbf{I}_{N \times N} & \mathbf{0}_{N \times N} \\ \mathbf{0}_{N \times N} & \mathbf{0}_{N \times N} \end{bmatrix}, \quad (17)$$

when noisy position observations are available. This nonetheless can be easily extended to other scenarios when for example velocities are observed. We will in practice assume that the noise parameters $\sigma_{1:N}$ and σ_v are pre-specified, e.g. by considering the sensor(s) characteristics.

After deducing the kinematic state for N objects in each dimension, under the assumption of dimension independence, the joint state for objects moving in a d-dimensional space can be easily expanded to $X_n = [X_n^1, \dots, X_n^d]$.

A simple state-independent leadership transition is proposed in this paper such that the sequence of L_n (i.e., leadership at the time instant t_n) follows a first-order discrete time Markov chain with the following transition probability

$$p(L_n | L_{n-1}, X_{n-1}) = p(L_n | L_{n-1}), \quad (18)$$

where $p(L_n | L_{n-1})$ denotes a fixed transition probability. For examples of state-dependent transition probabilities in groups see e.g., [10, 23]. Here, we specify a higher probability to an unchanged leadership pattern $L_n = L_{n-1}$, compared to $L_n \neq L_{n-1}$, as the leadership pattern is typically slowly changing and not subject to rapid switching; more details are provided in the experimental results in Section V-B. In our simulations, we adopt a switching prior in which all leadership structures may be switched to with equal probabilities (although the fish school real data example limits the number of leaders to 2). We apply such a generic scheme in order to illustrate the very effective exploration of the state space by our methods; however, more specific prior knowledge can easily be incorporated into the model, if needed, such as limiting the total number of leaders or permitting only certain objects to take leadership roles. This will enhance the efficiency of the estimation by reducing the numbers of possible leadership changes at the next time step. We note that it would also be possible within our framework to allow the state switching process to operate in continuous time, and this could be accomplished using for example a variable rate particle filter algorithm [11], although this is not explored in the current work.

III. STATE ESTIMATION

Given the sequence of observations from time step 0 to n , $Z_{0:n}$, a principal objective is to compute the joint probability density $p(X_n, L_{0:n}|Z_{0:n})$, which can be expanded as:

$$p(X_n, L_{0:n}|Z_{0:n}) = p(L_{0:n}|Z_{0:n})p(X_n|L_{0:n}, Z_{0:n}). \quad (19)$$

The posterior distribution $p(L_{0:n}|Z_{0:n})$ is not in a standard form and therefore sampling methods such as sequential Markov chain Monte Carlo (SMCMC) can be utilised to infer sequentially the leadership state, although the tractable Gaussian structure of the conditional model $p(X_n|L_{0:n}, Z_{0:n})$ leads to substantial simplifications. The marginal posterior distribution $p(L_{0:n}|Z_{0:n})$ is approximated by a set of unweighted particles $\{L_{0:n}^{(i)}\}_{i=1}^{N_p}$ obtained at time instant t_n as per:

$$\hat{p}(L_{0:n}|Z_{0:n}) = \frac{1}{N_p} \sum_{i=1}^{N_p} \delta_{L_{0:n}^{(i)}}(L_{0:n}), \quad (20)$$

where N_p denotes the number of the particles, $Z_{0:n}$ is the observation sequence up to time step n , and $\delta_y(x)$ is the Dirac function located at $x = y$. Detailed schemes for sequential sampling in this framework are presented in Section IV.

Conditional on the leadership state $L_{0:n}$, the dynamic model for X_n is linear and Gaussian. For this partially tractable state space model, a Rao-Blackwellisation scheme can be applied to improve the efficiency of the algorithm. In this way, the marginal distribution of the linear part X_n can be estimated from the leadership samples, resulting in a Gaussian mixture Monte Carlo approximation:

$$\begin{aligned} p(X_n|Z_{0:n}) &= \int p(X_n, L_{0:n}|Z_{0:n})dL_{0:n} \\ &= \int p(L_{0:n}|Z_{0:n})p(X_n|L_{0:n}, Z_{0:n})dL_{0:n} \\ &= \frac{1}{N_p} \sum_{i=1}^{N_p} \int p(X_n|L_{0:n}^{(i)}, Z_{0:n})\delta_{L_{0:n}^{(i)}}(L_{0:n})dL_{0:n} \\ &\approx \frac{1}{N_p} \sum_{i=1}^{N_p} p(X_n|L_{0:n}^{(i)}, Z_{0:n}), \end{aligned} \quad (21)$$

where every conditional distribution $p(X_n|L_{0:n}^{(i)}, Z_{0:n}) = \mathcal{N}(\mu_{n|0:n}^{(i)}, P_{n|0:n}^{(i)})$ may be calculated by Kalman filtering as follows:

$$\mu_{n|0:n-1}^{(i)} = F_{L_n}^{(i)} \mu_{n-1|0:n-1}^{(i)}, \quad (22)$$

$$P_{n|0:n-1}^{(i)} = F_{L_n}^{(i)} P_{n-1|0:n-1}^{(i)} F_{L_n}^{(i)T} + P_\tau, \quad (23)$$

$$\mu_{n|0:n}^{(i)} = \mu_{n|0:n-1}^{(i)} + k_n(Z_n - H\mu_{n|0:n-1}^{(i)}), \quad (24)$$

$$P_{n|0:n}^{(i)} = (I - k_n H) P_{n|0:n-1}^{(i)}, \quad (25)$$

$$k_n = P_{n|0:n-1}^{(i)} H^T (H P_{n|0:n-1}^{(i)} H^T + Q_v)^{-1}. \quad (26)$$

The likelihood is then $p(Z_n|Z_{0:n-1}, L_{0:n}^{(i)}) = \mathcal{N}(\mu_{Z_n}^{(i)}, P_{Z_n}^{(i)})$, whose mean and covariance are calculated by the prediction error decomposition [13]:

$$\begin{aligned} \mu_{Z_n}^{(i)} &= H\mu_{n|0:n-1}^{(i)}, \\ P_{Z_n}^{(i)} &= H P_{n|0:n-1}^{(i)} H^T + Q_v. \end{aligned} \quad (27)$$

IV. LEADERSHIP INFERENCE ALGORITHM

For the leadership state $L_{0:n}$, the prediction and update step for the posterior distribution $p(L_{0:n}|Z_{0:n})$ are

$$p(L_{0:n}|Z_{0:n-1}) = p(L_{0:n-1}|Z_{0:n-1})p(L_n|L_{n-1}) \quad (28)$$

$$p(L_{0:n}|Z_{0:n}) = \frac{p(Z_n|Z_{0:n-1}, L_{0:n})p(L_{0:n}|Z_{0:n-1})}{p(Z_n|Z_{0:n-1})}, \quad (29)$$

where $p(Z_n|Z_{0:n-1})$ is a normalization constant equal to

$$p(Z_n|Z_{0:n-1}) = \int p(Z_n|Z_{0:n-1}, L_{0:n})p(L_{0:n}|Z_{0:n-1})dL_{0:n}. \quad (30)$$

Although a closed form expression for $p(L_{0:n}|Z_{0:n})$ and the normalising constant are analytically available, the search for the optimal leadership sequences involves an intractable growth of permutations as time evolves, and hence sequential Monte Carlo methods such as the regular particle filter might be suitable for sequential inference of the leadership state sequences. However, when the number of objects in a group increases, the general particle filter methods suffer from a well known degeneracy problem [2]. Markov chain Monte Carlo (MCMC) techniques have been shown to be effective in high-dimensional cases and sequential MCMC approaches have been developed recently, which can be employed in an online inference setting [7, 27, 35]. Therefore, here we adopt the SMCMC method to infer the leadership state. We present several different versions of the SMCMC algorithm, and its special case, an online Gibbs sampler. The performance of all these methods is then demonstrated in Section V.

A. Sequential Markov Chain Monte Carlo

In the SMCMC method, the joint leadership posterior $p(L_{0:n}|Z_{0:n})$ is specified as the stationary distribution of a Markov chain at time step n . From equations (28)-(29), this posterior distribution is expressed as:

$$\begin{aligned} p(L_{0:n}|Z_{0:n}) &\propto p(Z_n|Z_{0:n-1}, L_{0:n})p(L_n|L_{n-1}) \\ &\quad \times p(L_{0:n-1}|Z_{0:n-1}). \end{aligned} \quad (31)$$

During the inference routine, once the stationary distribution $p(L_{0:n-1}|Z_{0:n-1})$ is attained at time step $n-1$, it can be approximated empirically by a set of N_p unweighted samples from the converged chain:

$$\hat{p}(L_{0:n-1}|Z_{0:n-1}) = \frac{1}{N_p} \sum_{i=1}^{N_p} \delta_{L_{0:n-1}^{(i)}}(L_{0:n-1}). \quad (32)$$

By using this approximation, the posterior distribution $p(L_{0:n}|Z_{0:n})$ can be obtained sequentially from the last as:

$$\begin{aligned} p(L_{0:n}|Z_{0:n}) &\propto \frac{1}{N_p} \sum_{i=1}^{N_p} p(Z_n|Z_{0:n-1}, L_{0:n-1}^{(i)}, L_n) \\ &\quad \times p(L_n|L_{n-1}^{(i)})\delta_{L_{0:n-1}^{(i)}}(L_{0:n-1}). \end{aligned} \quad (33)$$

In this section, we introduce two different joint proposals for this SMCMC sampler, which includes a prior dynamics-based proposal and an optimal kernel proposal. Both of these are developed as independence samplers in which the newly

proposed sample $L_{0:n}^*$ at iteration m does not depend on the sample $L_{0:n-1}^{m-1}$ at the last iteration $m-1$, i.e. $q(L_{0:n}^*|L_{0:n-1}^{m-1}) = q(L_{0:n}^*)$. The proposals generated are then either rejected or accepted according to a Metropolis-Hastings (M-H) acceptance ratio ρ . The general process of this sequential MCMC is given in Algorithm 1.

1) *Prior proposal*: A joint proposal $q(L_{0:n}|Z_{0:n})$ is designed where L_n and $L_{0:n-1}$ are updated jointly. Specifically, the proposal is expressed as:

$$q(L_{0:n}|Z_{0:n}) = \frac{1}{N_p} p(L_n|L_{n-1}) \sum_{i=1}^{N_p} \delta_{L_{0:n-1}^{(i)}}(L_{0:n-1}), \quad (34)$$

where we first sample $L_{0:n-1}^*$ from $\hat{p}(L_{0:n-1}|Z_{0:n-1})$ as in (32), and then propose L_n^* from $p(L_n|L_{n-1}^*)$. In this case, the acceptance function ρ' is calculated using:

$$\begin{aligned} \rho' &= \frac{p(L_n^*, L_{0:n-1}^*|Z_{0:n}) q(L_{0:n-1}^{m-1}|Z_{0:n})}{p(L_{n-1}^{m-1}, L_{0:n-1}^{m-1}|Z_{0:n}) q(L_{0:n}^*|Z_{0:n})} \\ &= \frac{p(Z_n|Z_{0:n-1}, L_{0:n-1}^*)}{p(Z_n|Z_{0:n-1}, L_{0:n-1}^{m-1})}. \end{aligned} \quad (35)$$

Consequently, the proposed samples are either accepted or rejected according to the acceptance ratio:

$$\rho = \min \left(1, \frac{p(Z_n|Z_{0:n-1}, L_{0:n-1}^*)}{p(Z_n|Z_{0:n-1}, L_{0:n-1}^{m-1})} \right). \quad (36)$$

After burn-in time N_{burn} , we store samples once every N_{thin} points and discard the others in order to reduce the autocorrelation. The posterior is then approximated by the resulting set of converged particles $\{L_{0:n}^{(i)}\}_{i=1}^{N_p}$. The filtering distribution $p(L_n|Z_{0:n})$ can also be easily attained by marginalising $L_{0:n-1}$ from the Monte Carlo representation, requiring no direct integration. This approach, however, entails a computational complexity of order $\mathcal{O}(N_p^2)$.

2) *Optimal proposal*: In the optimal situation, the joint proposal will be in the form of the approximated sought distribution at time n , i.e. $p(L_{0:n}|Z_{0:n})$,

$$\begin{aligned} q(L_{0:n}|Z_{0:n}) &= p(L_{0:n}|Z_{0:n}) \\ &= p(L_n|L_{0:n-1}, Z_{0:n}) \frac{p(L_{0:n-1}, Z_n|Z_{0:n-1})}{p(Z_n|Z_{0:n-1})} \\ &\propto p(L_n|L_{0:n-1}, Z_{0:n}) p(Z_n|L_{0:n-1}, Z_{0:n-1}) p(L_{0:n-1}|Z_{0:n-1}) \\ &= \frac{1}{N_p} p(L_n|L_{0:n-1}, Z_{0:n}) \sum_{i=1}^{N_p} p(Z_n|L_{0:n-1}^{(i)}, Z_{0:n-1}) \\ &\quad \times \delta_{L_{0:n-1}^{(i)}}(L_{0:n-1}) \end{aligned} \quad (37)$$

To obtain the joint samples, we first sample $L_{0:n-1}^*$ from $\frac{1}{N_p} \sum_{i=1}^{N_p} p(Z_n|L_{0:n-1}^{(i)}, Z_{0:n-1}) \delta_{L_{0:n-1}^{(i)}}(L_{0:n-1})$, and then propose L_n^* from $p(L_n|L_{0:n-1}^*, Z_{0:n})$.

In general, the challenge in utilising this optimal kernel is that we can neither sample from the proposal $p(L_n|L_{0:n-1}, Z_{0:n})$ nor calculate $p(Z_n|L_{0:n-1}, Z_{0:n-1})$, see [9] for the importance sampling case. In this paper, however, both of these distributions have closed form due to the discrete nature of leadership state and the linear Gaussian form for

the object dynamics. In particular, for each particle i , the likelihood $p(Z_n|L_{0:n-1}^{(i)}, Z_{0:n-1})$ can be evaluated as:

$$\begin{aligned} p(Z_n|L_{0:n-1}^{(i)}, Z_{0:n-1}) &= \int p(Z_n|L_{0:n-1}^{(i)}, L_n, Z_{0:n-1}) p(L_n = j|L_{0:n-1}^{(i)}, Z_{0:n-1}) dL_n \\ &= \frac{1}{K} \sum_{L_n \in S} p(Z_n|L_{0:n-1}^{(i)}, L_n, Z_{0:n-1}) p(L_n|L_{n-1}^{(i)}) \end{aligned} \quad (38)$$

where K is the number of possible leadership structures.

Similarly, the distribution $p(L_n|L_{0:n-1}^*, Z_{0:n})$ can be deduced as a discrete distribution:

$$\begin{aligned} p(L_n|Z_{0:n}, L_{0:n-1}^*) &= \frac{p(L_n, L_{0:n-1}^*|Z_{0:n})}{p(L_{0:n-1}^*|Z_{0:n})} \\ &\propto p(Z_n|Z_{0:n-1}, L_{0:n-1}^*, L_n) p(L_{0:n-1}^*|Z_{0:n-1}) p(L_n|L_{n-1}^*) \\ &= \frac{1}{K} \sum_{S_k \in S} p(Z_n|Z_{0:n-1}, L_{0:n-1}^*, L_n = S_k) \\ &\quad \times p(L_n = S_k|L_{0:n-1}^*) \delta_{S_k}(L_n), \end{aligned} \quad (39)$$

where $p(L_{0:n-1}^*|Z_{0:n-1})$ is a constant when $L_{0:n-1}^*$ is known. The value of likelihood $p(Z_n|Z_{0:n-1}, L_{0:n-1}^*, L_n)$ and the transition density $p(L_n|L_{n-1}^*)$ can be individually computed according to different values of $\{L_n = S_k\}_{k=1}^K$.

As the proposal is chosen as the stationary distribution, i.e., the leadership posterior, the acceptance function ρ here is calculated as

$$\rho = \frac{p(L_{0:n}^*|Z_{0:n}) q(L_{0:n-1}^{m-1})}{p(L_{0:n-1}^{m-1}|Z_{0:n}) q(L_{0:n}^*)} = 1 \quad (40)$$

Therefore, the joint proposed samples are always accepted.

In fact, since we are targeting the joint approximate posterior directly, we can see that this has reduced to an example of standard Monte Carlo and there is no MCMC burn-in required. The samples are still nonetheless approximate owing to the Monte Carlo approximation that is incorporated at each step as $\hat{p}(L_{0:n-1}|Z_{0:n-1})$.

Algorithm 1: General sequential MCMC algorithm

- 1 Initialization: $\{L_0^{(i)}\}_{i=1}^{N_p}$
 - 2 **for** time step $n = 0$ to T **do**
 - 3 initialization
 - 4 **for** iteration $m = 1$ to M **do**
 - 5 sample $L_{0:n}^*$ from $q(L_{0:n})$
 - 6 compute the acceptance ratio:

$$\rho = \min \left(1, \frac{p(L_{0:n}^*|Z_{0:n}) q(L_{0:n-1}^{m-1})}{p(L_{0:n-1}^{m-1}|Z_{0:n}) q(L_{0:n}^*)} \right)$$
 - 7 accept $(L_n^m, L_{0:n-1}^m) = (L_n^*, L_{0:n-1}^*)$
 - 8 After a burn-in time
 - 9 **for** $i = 1 : N_p$ **do**
 - 10 store $L_{0:n}^{(i)} = L_{0:n}^m$
 - 11 $p(L_{0:n}|Z_{0:n}) = \frac{1}{N_p} \sum_{i=1}^{N_p} \delta(L_{0:n} - L_{0:n}^{(i)})$
-

3) *Online Gibbs sampler*: The Gibbs sampler is a special case of the M-H algorithm in which the acceptance ratio is always one. In this part, we propose an online Gibbs sampler which can, as before, sequentially infer the leadership parameter $L_{0:n}$ with the joint distribution $p(L_{0:n}|Z_{0:n})$ as the stationary distribution at time step n . The premise of utilising the Gibbs sampler here is that both full conditionals $p(L_n|Z_{0:n}, L_{0:n-1})$ and $p(L_{0:n-1}|Z_{0:n}, L_n)$ have closed forms. The conditional distribution $p(L_n|Z_{0:n}, L_{0:n-1}^m)$ yields the following discrete distribution by using the approximation in (32):

$$\begin{aligned} \hat{p}(L_n|Z_{0:n}, L_{0:n-1}^m) &\propto p(Z_n|Z_{0:n-1}, L_{0:n-1}^m, L_n) \\ &\times p(L_n|L_{0:n-1}^m)p(L_{0:n-1}^m|Z_{0:n-1}) \\ &= \frac{1}{K} \sum_{S_k \in S} p(Z_n|Z_{0:n-1}, L_{0:n-1}^m, L_n = S_k) \\ &\times p(L_n = S_k|L_{0:n-1}^m)\delta_{S_k}(L_n). \end{aligned} \quad (41)$$

Having computed the sample L_n^m , the $L_{0:n-1}^m$ can be sampled from the condition distribution $p(L_{0:n-1}|Z_{0:n}, L_n^m)$ with its analytical expression in a form of a discrete distribution as per:

$$\begin{aligned} \hat{p}(L_{0:n-1}|Z_{0:n}, L_n^m) &\propto \frac{1}{N_p} \sum_{i=1}^{N_p} p(Z_n|Z_{0:n-1}, L_{0:n-1}^{(i)}, L_n^m) \\ &\times p(L_{0:n-1}^m|L_{0:n-1}^{(i)})\delta_{L_{0:n-1}^{(i)}}(L_{0:n-1}) \end{aligned} \quad (42)$$

where N_p is the total number of the samples. The process of Gibbs sampler is summarised in Algorithm 2.

Algorithm 2: Gibbs sampler

```

1 Initialization:  $\{L_0^{(i)}\}_{i=1}^{N_p}$ 
2 for time scan  $n = 0$  to  $T$  do
3   Initialization
4   for iteration  $m = 0$  to  $M - 1$  do
5     Given  $L_{n-1}^{m-1}$ , sample
6      $L_n^m \sim p(L_n^m|Z_{0:n}, L_{0:n-1}^{m-1})$  as in (41);
7     Given  $L_n^m$ , sample
8      $L_{0:n-1}^m \sim p(L_{0:n-1}^m|Z_{0:n}, L_n^m)$  as in (42);
8   Repeat the process for  $M$  times. After a burn-in
   time, keep the subsequent  $N_p$  converged samples
    $\{L_{n-1}^{(i)}, L_n^{(i)}\}_{i=1}^{N_p}$ .
```

Compared to the SMCMC-prior sampler with a computational complexity of $\mathcal{O}(N_p)$, both the SMCMC-optimal and online Gibbs sample have a computation complexity of $\mathcal{O}(N_p^2)$.

B. Marginal Particle Filter

As the state of leadership parameter is discrete and finite, we could alternatively implement the marginal particle filter [15][12]. As before, the Rao-blackwellisation is applied and the kinematic state distribution of objects are estimated by using Kalman filtering.

In the marginal particle filter, the filtering distribution $p(L_n|Z_{0:n})$ is utilised as the target distribution instead of

the joint posterior $p(L_{0:n}|Z_{0:n})$, with purpose of mitigating the degeneration problem of particle filter in high-dimension case. Assume the posterior distribution $p(L_{0:n-1}|Z_{0:n-1})$ is approximated empirically by a set of weighted particles $\{L_{n-1}^{(i)}, w_{n-1}^{(i)}\}_{i=1}^{N_p}$ obtained at time step $n - 1$:

$$\hat{p}(L_{0:n-1}|Z_{0:n-1}) = \frac{1}{N_p} \sum_{i=1}^{N_p} w_{n-1}^{(i)} \delta_{L_{0:n-1}^{(i)}}(L_{0:n-1}), \quad (43)$$

where the weight $\sum_{i=1}^{N_p} w_{n-1}^{(i)} = 1$.

It follows that the prediction step can be formulated as:

$$p(L_n|Z_{0:n-1}) = \frac{1}{N_p} \sum_{j=1}^{N_p} w_{n-1}^{(j)} p(L_n|L_{n-1}^{(j)}), \quad (44)$$

and the update step at time n is written as:

$$p(L_n|Z_{0:n}) \propto \sum_{j=1}^{N_p} w_{n-1}^{(j)} p(Z_n|Z_{0:n-1}, L_{0:n-1}^{(j)}, L_n) p(L_n|L_{n-1}^{(j)}). \quad (45)$$

The general process of this marginal particle filter is given in Algorithm 3. At each time step n , each particle $L_n^{(i)}$ is first sampled from the defined proposal $q(L_n|Z_{0:n})$, which assumes to have the form of:

$$q(L_n|Z_{0:n}) = \sum_{j=1}^{N_p} w_{n-1}^{(j)} q(L_n|L_{0:n-1}^{(j)}, Z_{0:n}). \quad (46)$$

Subsequently, the corresponding weight is calculated via:

$$w_n^{(i)} = \frac{p(L_n^{(i)}|Z_{0:n})}{q(L_n^{(i)}|Z_{0:n})}, \quad (47)$$

and it is further normalised to make sure the sum of the weights equals one.

Similarly, here we analyse two different proposals, the prior and optimal proposal, as comparisons to the above SMCMC methods.

Algorithm 3: General marginal particle filter

```

1 Initialization:  $\{L_0^{(i)}, w_0^{(i)}\}_{i=1}^{N_p}$ 
2 for time scan  $n = 1$  to  $T$  do
3   for particle  $i = 1$  to  $N_p$  do
4     sample  $L_n^{(i)}$  from  $q(L_n|Z_{0:n})$  as in (46)
5     evaluate the weight  $\hat{w}_n^{(i)}$  as in (47):
6   Normalization:  $w_n^{(i)} = \frac{\hat{w}_n^{(i)}}{\sum_{j=1}^{N_p} \hat{w}_n^{(j)}}$ 
7    $p(L_n|Z_{0:n}) = \frac{1}{N_p} \sum_{i=1}^{N_p} w_n^{(i)} \delta(L_n - L_n^{(i)})$ 
```

1) *Prior proposal*: Specifically, the prior proposal to generate L_n in Algorithm 3 is designed as:

$$q(L_n|Z_{0:n}) = \sum_{j=1}^{N_p} w_{n-1}^{(j)} p(L_n|L_{n-1}^{(j)}). \quad (48)$$

After we have sampled $L_n^{(i)}$, the weight is computed as per:

$$w_n^{(i)} = \frac{\sum_{j=1}^{N_p} w_{n-1}^{(j)} p(Z_n | Z_{0:n-1}, L_{0:n-1}^{(j)}, L_n^{(i)}) p(L_n^{(i)} | L_{n-1}^{(j)})}{\sum_{j=1}^{N_p} w_{n-1}^{(j)} p(L_n^{(i)} | L_{n-1}^{(j)})}. \quad (49)$$

The general step of marginal particle filter using prior proposal endures a computational complexity of $\mathcal{O}(N_p^2)$, whereas that of the standard particle filters is $\mathcal{O}(N_p)$. Techniques have been proposed to reduce the computation complexity to $\mathcal{O}(N_p \log N_p)$, see [15] for more details.

2) *Optimal proposal*: In this approach, we assume the optimal proposal:

$$q(L_n | Z_{0:n}) = \sum_{j=1}^{N_p} w_{n-1}^{(j)} p(L_n | L_{0:n-1}^{(j)}, Z_{0:n}). \quad (50)$$

After we have sampled $L_n^{(i)}$, the weight for the particle i is computed according to:

$$\begin{aligned} w_n^{(i)} &= \frac{\sum_{j=1}^{N_p} w_{n-1}^{(j)} p(Z_n | Z_{0:n-1}, L_{0:n-1}^{(j)}, L_n^{(i)}) p(L_n^{(i)} | L_{0:n-1}^{(j)}, Z_{0:n})}{\sum_{j=1}^{N_p} w_{n-1}^{(j)} p(L_n^{(i)} | L_{0:n-1}^{(j)}, Z_{0:n})} \\ &= p(Z_n | Z_{0:n-1}, L_{0:n-1}^{(i)}) \end{aligned} \quad (51)$$

We can observe that the samples drawn are equivalent marginally to the draws from the joint optimal proposal in Section IV-A. Here, this scheme again can be seen to reduce to a standard Monte Carlo.

V. RESULT

In this section, we evaluate the proposed dynamic leadership inference approach on both synthetic and real animal data. For the former, two datasets with different parameter settings are generated in order to compare the leadership state estimation performance of the different sequential Monte Carlo algorithms introduced in Section IV. Subsequently, we assess the performance of the developed techniques for the task of analysing animal collective behaviour. In particular, the dynamic leadership pattern in two moving animal groups is examined; they are a homing pigeon flock [21] and a fish school [29].

A. Synthetic Data

The two synthetic datasets used here are simulated utilising the dynamic model in Section II, and each dataset includes 100 Monte Carlo runs. For dataset 1, we have observations of a group of four objects moving over 100 time steps, and the time interval between observations is $\tau = 1$ s. The group model parameters are: $\alpha = 0.2$, $\beta = 0.2$, $\gamma_{1:N} = 0.1$ and $\eta = 0.005$. Whereas, the noise parameters are set as $\sigma_{1:N} = 2$ and $\sigma_v = 1$. The destination D and the leadership pattern $L_{0:T}$ are assigned randomly over 100 Monte Carlo. In dataset 2, we double the group size in order to evaluate the performance of these algorithms in a higher dimension case and observations of eight objects are generated; all other parameters are identical to those of dataset 1.

For all the sequential Monte Carlo methods we presented here, 1000 particles are employed to approximate $p(L_{0:n} | Z_{0:n})$

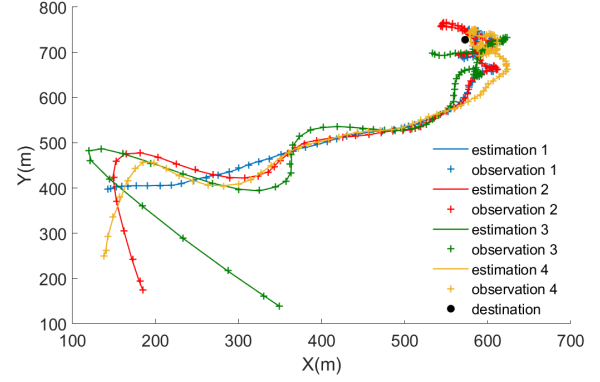


Fig. 2: Example synthetic tracks with time-varying leadership; Black dot denotes the predefined destination. solid lines are the estimated tracks and pluses are observations.




Fig. 3: Estimated probability of leadership $p(L_n | Z_{0:n})$; black dotted line indicates the true leader L_n at each time step n .

at each time step n . The possible leadership structures of a group of N objects are defined to include all subsets of the set $\{1 \dots N\}$, except for itself and the empty set. Hence, the total number of the leadership structures will be $2^N - 2$. For example, for dataset 1, all possible leadership structures are listed in Table. I and the total number is 14. For dataset 2, the total number of possible leadership structures increases to 254, and therefore would be valuable to assess the outcomes of the applied inference algorithms under such conditions.

S_1	S_2	S_3	S_4	S_5	S_6	S_7
1	2	3	4	{1,2}	{1,3}	{1,4}
S_8	S_9	S_{10}	S_{11}	S_{12}	S_{13}	S_{14}
{2,3}	{2,4}	{3,4}	{1,2,3}	{1,2,4}	{1,3,4}	{2,3,4}

TABLE I: The list of leadership structures for dataset 1.

To visualise the performance of tracking and leadership inference, we analyse one example track from dataset 1. The observations and the estimated tracks of the four objects in the group are shown in Fig. 2. The ground-truth and estimated leadership probabilities $p(L_n | Z_{0:n})$ from the optimal kernel SMC sampler are depicted in Fig. 3. It clearly depicts the ability of the proposed approach to effectively capture the dynamic change of leadership pattern in a group over time.

Next, the average performance of SMC algorithms (using the prior and optimal proposals), Gibbs sampler and marginal particle filter (with prior and optimal proposals) are presented. The metric utilised for measuring the leadership in-

ference performance is the mean of the "correct rate" averaged over 100 runs. The correct rate is calculated by using the MAP principle as per: at each time step, we consider the inferred leadership with the highest probability as the current output leader, and then compare it with the ground truth; if they are equal, we assume to obtain one correct estimation. For each run, the final correct rate over 100 time steps is calculated as the total number of correct estimations divided by 100. The tracking performance is evaluated using the Root Means Square Error (RMSE) averaged over the 100 independent runs. Meanwhile, the computational complexity is also studied by calculating the CPU time (System: Intel(R) Core(TM) i7-8550 CPU@1.80 GHz, 8 GB RAM) required at one time step in each dataset and averaged over 100 Monte Carlo runs.

For dataset 1, Table II presents the mean of the correct leader estimation rates of the five tested techniques averaged across the 100 runs. It can be seen that all these examined methods can reliably estimate the leadership state. Although the difference is subtle, we can still observe that using the optimal kernel, in both SMC MC and marginal particle filter, deliver better results compared with using the prior proposal. Overall, SMC MC methods outperform marginal particle filter (except for the optimal kernel case, since this is statistically equivalent in both cases, even though the algorithmic implementations are different, see Section IV). For the execution times, all examined methods in the case of a four-object group size can potentially output leadership inference results in real-time with the testing system (assuming that the recorded execution time is less than the duration between two successive sensory observations). SMC MC using prior proposal has the shortest execution time (i.e. lowest computational complexity) and Gibbs the longest. Additionally, the mean of RMSE over 100 Monte Carlo runs are displayed in Table II. For models that include the leadership pattern, we can see that all methods show promising, nearly similar, tracking performance as we adopt the Rao-Blackswellisation scheme. The marginal PF with prior proposal exhibits a slight less accurate tracking results due to its comparatively poorer estimation of the true leadership state. In contrast, the model without considering leadership structure has noticeably higher tracking RMSE. This demonstrate that incorporating interaction among group members enhances the achieved tracking accuracy.

For dataset 2, the overall performance is shown in Table III. It can be noticed that the algorithms using the optimal proposal have the highest correct rate. With the prior proposal, the correct rate of SMC MC is substantially higher than that delivered by the marginal particle filter. This illustrates the SMC MC superiority in high dimension cases, i.e. more scalable, and therefore we can conclude that the SMC MC methods outperform the marginal particle filters. The run-time grows markedly with the number of group members for all the algorithms, with SMC MC using the prior proposal having a significantly shorter execution time. This observation is consistent with the analysis of the computational complexity in Section IV where the SMC MC-prior method has a computational complexity of $\mathcal{O}(N_p)$ compared with $\mathcal{O}(N_p^2)$ for all the other techniques. Model without considering the leadership pattern again has the largest mean of the tracking RMSE over

the 100 Monte Carlo runs.

TABLE II: Overall performance for dataset 1 over 100 runs.

	Gibbs	SMC MC		Marginal PF		without leadership
		Prior	Optimal	Prior	Optimal	
Correct rate	0.87	0.85	0.87	0.79	0.87	-
Time (sec)	0.67	0.14	0.31	0.25	0.31	-
RMSE	0.27	0.27	0.27	0.30	0.27	0.67

TABLE III: Overall performance for dataset 2 over 100 runs.

	Gibbs	SMC MC		Marginal PF		without leadership
		Prior	Optimal	Prior	Optimal	
Correct rate	0.78	0.76	0.79	0.65	0.79	-
Time (sec)	62.8	1.39	45.1	57.2	45.1	-
RMSE	0.20	0.20	0.20	0.21	0.20	0.35

B. Real Data

In this part, we only present the result from the proposed SMC MC sampler with optimal kernel to two real animal datasets, since the methods using optimal sampler have shown a better performance with the synthetic data. It is noted that the MPF sampler with the optimal kernel is statistically equivalent to the SMC MC sampler with optimal kernel in our case.

TABLE IV: Overall tracking performance for two real scenarios.

Dataset	RMSE	
	with leadership	without leadership
Pigeon flock	3.65	4.07
Fish school	2.71	2.96

1) *Homing Pigeon Flock*: Here we analyse a 1000-time-step-segment flight of a flock of four homing pigeons during a journey, reconstructed from the available high-precision GPS data collected in [21]. The time interval between observations is 0.2s. The observations and estimated tracks of the four pigeons are given in Fig. 4. It can be seen that the flock trajectories are estimated accurately. Below, we compare our methods that consider the leadership structure with the standard CV model without any group interactions. The RMSE results are shown in Table IV. It can be seen that considering the leadership structure enhances the overall tracking performance.

The estimated leadership of the flock data over 1000 time steps is depicted in Fig. 5. The color bar ranging from white to red indicates the probability of this pigeon being the leader deduced at each time step. We can see that, despite assuming that all the 14 leadership structures in S are possible, the most probable leader in the homing pigeon flock is inferred as object 1 for the first 920 time steps; after the object 1 arrives at the destination around time step 920, the object 2 takes the lead for the following 80 time steps, leading the other pigeons to their intended destination (i.e. home). Such result is consistent with the dominance hierarchy studied in [21] about this pigeon flock. Different from the method in [21], which determines the

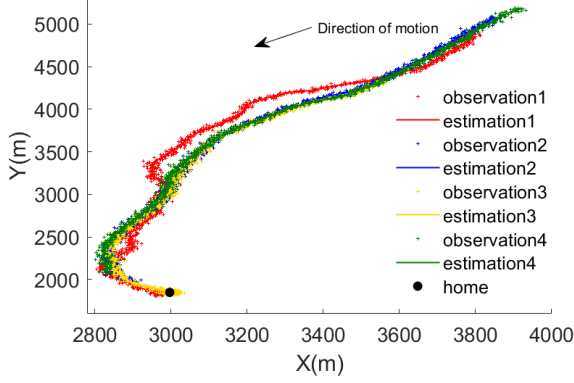


Fig. 4: Estimated Pigeons tracks; solid lines are estimations, and pluses are the available observations. The black dot denotes the home of pigeon flock




Fig. 5: The estimated probability of leadership in the homing pigeon flock.

leadership hierarchy by direct calculation of the directional correlation delay times, here we reveal the leadership pattern by using Bayesian inference methods.

2) *Fish school foraging scenario*: The performance of the introduced approach is further evaluated in a 50-time-step video clip of a school of five golden shiners fish in [29]. The video is recorded by a Sony EX1 video camera (1280 x 720 pixels) mounted above the fish tank at a frame rate of 60 Hz, and then be pre-processed into a 2-dimensional position data. A comparatively-small part of the fish are trained to possess foraging ability; in other words, the trained fish will know the location of the feeder and recognise the signal that there will be food once the feeder's light is on. The other fish in the group are not informed of this foraging information. After the signal of food is given, there will be around 10-time-step response time for the trained fish to notice the light, and they will then dash to the location of the feeder (destination). Meanwhile, the other fish will follow the leaders to the intended endpoint (food location). The video clip recorded here starts before any visible response from the fish school is observed, and ends when the first fish reached the feeder. It is noticeable that even the fish have been trained, we still cannot confirm whether they will successfully play the leadership role as the fish might be ill-trained or for other reasons (e.g., the untrained fish cannot follow the trained fish due to occlusion in vision).

In this experiment, both fish 1 and fish 2 are trained before the experiment began, and this information is utilised as ground truth for evaluating the effectiveness of our method.




Fig. 6: Estimated tracks of a fish school; solid lines are the estimations, and pluses are observations. The black dot denotes the feeder in the fish tank




Fig. 7: The estimated probability of leadership in the fish school.

In performing tracking the leadership model is unaware that fish 1 and fish 2 are leaders and the algorithm attempts to determine this; however, the total number of trained fish in the group is known *a priori* to be two. This prior information implies that there are at most two leaders in this fish school. Thus, the list of possible leadership structures is shortened from 30 to 15, as displayed in Table V.

S_1	S_2	S_3	S_4	S_5
1	2	3	4	5
S_6	S_7	S_8	S_9	S_{10}
{1,2}	{1,3}	{1,4}	{1,5}	{2,3}
S_{11}	S_{12}	S_{13}	S_{14}	S_{15}
{2,4}	{2,5}	{3,4}	{3,5}	{4,5}

TABLE V: The list of possible leadership structures

The observations and estimated tracks of this fish school over 50 time steps are shown in Fig. 6. The black circle denotes the feeder's location, i.e., the destination. It can be seen that the proposed inference method can accurately estimate the kinematic states of the fish school. Similarly, to show the effectiveness of considering the leadership structure, we compare our methods with standard CV model, and the RMSE results are shown in Table IV. It can be seen that our proposed model has a smaller RMSE value, which demonstrates its enhanced tracking performance. From Fig. 6, we can also observe that around response time step 10, fast manoeuvres happen where there are abrupt changes in the fish velocities. The estimated probability of leadership in the fish school over

50 time step is depicted in Fig. 7. We can see that before time step 10, it provides little useful information about leadership pattern in the fish school, which corresponds to the respond time discussed before. After time step 10, fish 2 becomes the most probable leader of the group, with the leadership pattern changing to S_6 between time step 16 and 20 when fish 1 and 2 lead the group together. As the ground truth is that fish 1 and 2 are the informed/trained fish, we can see that the inferred leadership pattern is reasonable, which proves that our proposed method is capable of capturing the dynamic dominance hierarchy amongst the fish school.

VI. CONCLUSION

In this paper, we introduce a novel leader-follower model to capture the dynamic interactions amongst groups of objects. Compared with previous models, the leadership structure within the proposed formulation is significantly more flexible in that not only the group leadership can change over time, but also the number of group leaders. Additionally, we generalize our behaviour model for each leader such that its motion can be driven by reaching a final destination, which is commonly in numerous applications and real-life scenarios. Under the Bayesian inference framework, the group leadership state is inferred online by a sequential Markov chain Monte Carlo (SMCMC) algorithm, which is evaluated under different proposals and also compared with the well-known particle filter methods. The results presented demonstrate the effectiveness of the proposed Bayesian inference framework, and its potential applications in facilitating the understanding of animal group behaviours. The work has various extensions including the estimation of the group parameters, consideration of more complex hierarchical structures in larger groups/networks, and treating the data association step for group tracking scenarios when the objects and measurements assignments are unknown.

REFERENCES

- [1] Ahmad, B. I., Murphy, J. K., Langdon, P. M., Godsill, S. J., Hardy, R. and Skrypchuk, L. [2015], 'Intent inference for hand pointing gesture-based interactions in vehicles', *IEEE transactions on cybernetics* **46**(4), 878–889.
- [2] Bengtsson, T., Bickel, P., Li, B. et al. [2008], Curse-of-dimensionality revisited: Collapse of the particle filter in very large scale systems, in 'Probability and statistics: Essays in honor of David A. Freedman', Institute of Mathematical Statistics, pp. 316–334.
- [3] Carmi, A. Y., Mihaylova, L., Gning, A., Gurfil, P. and Godsill, S. [2013], Monte Carlo-based Bayesian group object tracking and causal reasoning, in 'Advances in Intelligent Signal Processing and Data Mining', Springer, pp. 7–53.
- [4] Carmi, A. Y., Mihaylova, L., Septier, F., Pang, S. K., Gurfil, P. and Godsill, S. [2011], MCMC-based tracking and identification of leaders in groups, in '2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)', IEEE, pp. 112–119.
- [5] Consolini, L., Morbidi, F., Prattichizzo, D. and Tosques, M. [2008], 'Leader-follower formation control of non-holonomic mobile robots with input constraints', *Automatica* **44**(5), 1343–1349.
- [6] Contractor, N. S., DeChurch, L. A., Carson, J., Carter, D. R. and Keegan, B. [2012], 'The topology of collective leadership', *The Leadership Quarterly* **23**(6), 994–1011.
- [7] De Freitas, A., Septier, F. and Mihaylova, L. [2015], 'Sequential Markov Chain Monte Carlo for Bayesian filtering with massive data', *arXiv preprint arXiv:1512.02452*.
- [8] Dong, L., Chai, S., Zhang, B., Nguang, S. K. and Wen, L. [2015], 'High-order tracking problem with a time-varying topology and communication delays', *Neurocomputing* **149**, 1360–1369.
- [9] Doucet, A., Godsill, S. and Andrieu, C. [2000], 'On sequential Monte Carlo sampling methods for Bayesian filtering', *Statistics and computing* **10**(3), 197–208.
- [10] Gning, A., Mihaylova, L., Maskell, S., Pang, S. K. and Godsill, S. [2010], 'Group object structure and state estimation with evolving networks and Monte Carlo methods', *IEEE Transactions on Signal Processing* **59**(4), 1383–1396.
- [11] Godsill, S. J., Vermaak, J., Ng, W. and Li, J. F. [2007], 'Models and algorithms for tracking of maneuvering objects using variable rate particle filters', *Proceedings of the IEEE* **95**(5), 925–952.
- [12] Gustafsson, F. [2013], On marginal particle filters with linear complexity, in '2013 5th IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)', IEEE, pp. 356–359.
- [13] Harvey, A. C. [1990], *Forecasting, structural time series models and the Kalman filter*, Cambridge university press.
- [14] Khan, Z., Balch, T. and Dellaert, F. [2004], An MCMC-based particle filter for tracking multiple interacting targets, in 'European Conference on Computer Vision', Springer, pp. 279–290.
- [15] Klaas, M., de Frietas, N. and Doucet, A. [2005], 'Toward a practical N2 Monte Carlo filter: the marginal particle filter', *Proc of Artificial Intelligence*.
- [16] Li, Q. and Godsill, S. [2020], A new leader-follower model for bayesian tracking, in '2020 IEEE 23rd International Conference on Information Fusion (FUSION)', IEEE, pp. 1–7.
- [17] Li, Q., Godsill, S. J., Liang, J. and Ahmad, B. I. [2020], Inferring dynamic group leadership using sequential Bayesian methods, in 'ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)', IEEE, pp. 8911–8915.
- [18] Mahler, R. P. [2007], *Statistical multisource-multitarget information fusion*, Vol. 685, Artech House Norwood, MA.
- [19] Mihaylova, L., Carmi, A. Y., Septier, F., Gning, A., Pang, S. K. and Godsill, S. [2014], 'Overview of Bayesian sequential Monte Carlo methods for group and extended object tracking', *Digital Signal Processing* **25**, 1–16.
- [20] Murphy, J., Ozkan, E., Bunch, P. and Godsill, S. [2016],

Sparse structure inference for group and network tracking, in '2016 19th International Conference on Information Fusion (FUSION)', IEEE, pp. 1208–1214.

- [21] Nagy, M., Vásárhelyi, G., Pettit, B., Mariani, I. R., Vicsek, T. and Biro, D. [2013], 'Context-dependent hierarchies in pigeons', *Proceedings of the National Academy of Sciences* **110**(32), 13049–13054.
- [22] Oh, S., Russell, S. and Sastry, S. [2009], 'Markov Chain Monte Carlo data association for multi-target tracking', *IEEE Transactions on Automatic Control* **54**(3), 481–497.
- [23] Pang, S. K., Li, J. and Godsill, S. [2011], 'Detection and tracking of coordinated groups', *IEEE Transactions on Aerospace and Electronic Systems* **47**(1), 472–502.
- [24] Parrish, J. K. and Hamner, W. M. [1997], *Animal groups in three dimensions: how species aggregate*, Cambridge University Press.
- [25] Reynolds, C. W. [1987], Flocks, herds and schools: A distributed behavioral model, in 'Proceedings of the 14th annual conference on Computer graphics and interactive techniques', pp. 25–34.
- [26] Ristic, B., Arulampalam, S. and Gordon, N. [2003], *Beyond the Kalman filter: Particle filters for tracking applications*, Artech house.
- [27] Septier, F., Pang, S. K., Carmi, A. and Godsill, S. [2009], On MCMC-based particle methods for Bayesian filtering: Application to multitarget tracking, in '2009 3rd IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)', IEEE, pp. 360–363.
- [28] Shrit, O., Martin, S., Alagha, K. and Pujolle, G. [2017], A new approach to realize drone swarm using ad-hoc network, in '2017 16th Annual Mediterranean Ad Hoc Networking Workshop (Med-Hoc-Net)', IEEE, pp. 1–5.
- [29] Strandburg-Peshkin, A., Twomey, C. R., Bode, N. W., Kao, A. B., Katz, Y., Ioannou, C. C., Rosenthal, S. B., Torney, C. J., Wu, H. S., Levin, S. A. et al. [2013], 'Visual sensory networks and effective information transfer in animal groups', *Current Biology* **23**(17), R709–R711.
- [30] Särkkä, S. and A. Solin [2019], *Applied Stochastic Differential Equations*, Institute of Mathematical Statistics Textbooks, Cambridge University Press.
- [31] Tahir, A., Böling, J., Haghighyan, M.-H., Toivonen, H. T. and Plosila, J. [2019], 'Swarms of unmanned aerial vehicles—A survey', *Journal of Industrial Information Integration* **16**, 100106.
- [32] Vermaak, J., Godsill, S. J. and Perez, P. [2005], 'Monte Carlo filtering for multi target tracking and data association', *IEEE Transactions on Aerospace and Electronic Systems* **41**(1), 309–332.
- [33] Von Stackelberg, H. [2010], *Market structure and equilibrium*, Springer Science & Business Media.
- [34] Xing, J., Ai, H., Liu, L., Lao, S. et al. [2011], 'Multiple player tracking in sports video: A dual-mode two-way Bayesian inference approach with progressive observation modeling'.
- [35] Yang, Y. and Dunson, D. B. [2013], 'Sequential Markov Chain Monte Carlo', *arXiv preprint arXiv:1308.3861*.



Qing Li received the B.E. degree in information engineering from the Nanjing University of Aeronautics and Astronautics, China, in 2015, and the M.E. degree in information and communication engineering from Beihang University, China, in 2018. She is a Ph.D. student in the Signal Processing and Communications Laboratory at the Department of Engineering, Cambridge University, U.K. since 2018. Her research interests include Monte Carlo computational methods, Bayesian statistical inference, and tracking.



Bashar I. Ahmad, is a signal processing specialist at Aveillant, THALES Land and Air Systems (LAS) since January 2019 and a senior research associate at the Engineering Department, Cambridge University since 2013; he was elected a Research Fellow of Wolfson College, Cambridge. Prior to joining Cambridge University, he was a postdoctoral research associate at Imperial College London, U.K., a project manager at DSPG Telecom (BT), U.K., and a research engineer at Philips/NXP Semiconductors, U.K. Bashar received the B.Eng. degree (Hons.) in electronic engineering and the Ph.D. degree from University of Westminster, U.K., in 2007 and 2011, respectively; he also undertook the Msc in Signal Processing and Communications (selected modules) at Imperial College London, U.K. His research areas of interest include the theory and applications of statistical signal processing and Bayesian inference for time-series data, such as in object tracking and sensor data fusion, automatic target recognition, automation, sub-Nyquist data acquisition systems, intelligent vehicles and others.



Simon J. Godsill Simon Godsill is Professor of Statistical Signal Processing in the Engineering Department at Cambridge University. He is also a Professorial Fellow and tutor at Corpus Christi College Cambridge. He coordinates an active research group in Signal Inference and its Applications and is Head of the Signal Processing and Communications (SigProC) Laboratory at Cambridge. His group specialises in Bayesian computational methodology, multiple object tracking, audio and music processing, and financial time series modeling. A particular methodological theme over recent years has been the development of novel techniques for optimal Bayesian filtering and smoothing, using Sequential Monte Carlo or Particle Filtering methods. Prof. Godsill has published extensively in journals, books and international conference proceedings, and has given a number of high profile invited and plenary addresses at conferences such as the Valencia conference on Bayesian Statistics (twice), the IEEE Statistical Signal Processing Workshop, the Conference on Bayesian Inference for Stochastic Processes (BISP), the IEEE Workshop on Machine Learning in Signal Processing (2013) and FUSION (2016). He co-authored a Springer text Digital Audio Restoration with Prof. Peter Rayner in 1998. He was technical chair of the IEEE NSSPW workshop in 2006 on sequential and nonlinear filtering methods, and has been on the conference panel for numerous other conferences/workshops. Prof. Godsill has served as Associate Editor for IEEE Tr. Signal Processing and the journal Bayesian Analysis. He was Theme Leader in Tracking and Reasoning over Time for the UK's Data and Information Fusion Defence Technology Centre (DIF-DTC) and Principal Investigator on many grants funded by the EU, EPSRC, QinetiQ, General Dynamics, MOD, Microsoft UK, Citibank, Mastercard, Google, DSO Singapore, Huawei and Jaguar Landrover. In 2009-10 he was co-organiser of an 18 month research program in Sequential Monte Carlo Methods at the SAMSI Institute in North Carolina and in 2014 he co-organised a research programme at the Isaac Newton Institute on Sequential Monte Carlo methods. In 2018 he will be General Chair of the FUSION Conference in Cambridge. Two of his journal papers have recently received Best Paper awards from the IEEE and IET. He continues to be a Director of CEDAR Audio Ltd. (which has received numerous accolades over the years, including a technical Oscar), and for which he was a founding staff member in 1988. The company has commercialised many of the ideas from Professor Godsill's research over the years.