

Multi-Model Motion Tracking under Multiple Team Member Actuators

Yang Gu
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213, USA
guyang@cs.cmu.edu

Manuela Veloso
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213, USA
mmv@cs.cmu.edu

ABSTRACT

Autonomous robots need to track objects. Object tracking relies on predefined robot motion and sensory models. Tracking is particularly challenging if the robots can actuate on the object to be tracked, as the motion can become highly discontinuous and nonlinear. We have previously developed a successful tracking approach that switches among target motion models as a function of one robot's actions. In this paper, we consider the object to be effected by a team of agents. We contribute on our team-based tracking method that can use a dynamic multi-motion model based on a team coordination plan. We present the multi-target multi-model probabilistic tracking algorithm in detail and present empirical results both in simulation and in a human-robot Segway soccer team. The team coordination plan allows the robot to much more effectively track mobile targets.

Categories and Subject Descriptors

I.2.9 [Artificial Intelligence]: Robotics—*Sensors, Autonomous vehicles, Kinematics and dynamics*

General Terms

Algorithms, Experimentation, Human Factors

Keywords

multi-model, motion modeling, tracking, actuator

1. INTRODUCTION

Autonomous robot agents need to be able to track moving objects. (e.g. [11]). When tracking is performed by a robot executing specific tasks acting over the target being tracked, such as a Segway RMP soccer robot grabbing and kicking a ball, the motion model of the target becomes dependent on the robot's actions [9]. The robot's tactic provides valuable information in terms of the target behavior. A tactic-based

motion modeling and tracking in such scenarios has been introduced in [7].

However, the Segway RMP soccer robot operates in such an environment that all the players on the field can actuate over the ball, namely grab and kick the ball according to the rules, which makes the motion model of the ball even more complex. When the robot is joining a game as a member of a human-robot team, the team coordination knowledge provides further information that can be incorporated into the motion modeling and tracking process. Furthermore, the approach in [7] does not perform well once incorrect measurements originating from clutter or false alarms exist, which causes multiple hypothesis of the tracked target. Recently, a hybrid approach for online joint detection and tracking for multiple targets was proposed [10]. This approach does not assume the knowledge of true targets (without clutter) is given. It first uses a deterministic clustering method that searches for regions of interest (ROIs) based on the observations and monitors these ROIs for target detection, then performs multi-target tracking by Sequential Monte Carlo (SMC) methods. In this paper, based on their approach, we present an extension to the tactic-based tracking scheme introduced in [7] to solve a plan-dependent multi-target multi-hypothesis tracking problem.

The paper is organized as follows. We first give a brief description of the Segway RMP soccer robot. Next we show the system dynamics. We introduce play-based motion model which incorporates the team coordination knowledge. We then describe the multi-target multi-model tracking algorithm, leading to our experimental results, related work, and conclusions.

2. SEGWAY RMP SOCCER ROBOT

The Segway platform is unique due to its combination of wheel actuators and dynamic balancing [5]. Segway RMP, or Robot Mobility Platform, provides an extensible control platform for robotics research. It imbues the robot with the novel characteristics of a fast platform and travel long ranges, able to carry significant payloads, able to navigate in relatively tight spaces for its size, and provides the opportunity to mount sensors at a height comparable to human eye level [12].

In our previous work, we have developed a Segway RMP robot base capable of playing Segway soccer. We briefly describe the two major components of the control architecture, the sensor and the robot cognition, which are highly related to our multi-model motion tracking.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'06 May 8–12 2006, Hakodate, Hokkaido, Japan.
Copyright 2006 ACM 1-59593-303-4/06/0005 ...\$5.00.

2.1 Vision Sensor and Infrared Sensors

The goal of vision is to provide as many valid estimates of targets as possible. Tracking then fuses this information to track the most interesting targets (the ball and the team member, in this paper) of relevance to the robot. We use one pan-tilt camera as the vision sensor. We do not discuss the localization of the robot in the sense that a lot of soccer tasks (known as tactics and plays described in section 2.2) can be done by the Segway RMP robot without localization knowledge. Also we use global reference for position and velocity in this paper which means it is relative to the reference point where the robot starts to do dead reckoning.

We have equipped each robot with infrared sensors to reliably detect the objects located in the catchable area of the robot. Its measurement is a binary value indicating whether or not an object is there. In most cases, this is the blind area of the pan-tilt camera. Therefore, the infrared sensor is particularly useful when the robot is grabbing the ball.

2.2 Robot Cognition

A control architecture, called Skills-Tactics-Plays, was proposed in [4] to achieve the goals of responsive, adversarial team control. The key component of STP is the division between single robot behavior and team behavior.

A play, P , is a fixed team plan consisting of a set of applicability conditions, termination conditions, and N roles, one for each team member. Each role defines a sequence of tactics T_1, T_2, \dots and associated parameters to be performed by that role in the ordered sequence. Assignment of roles to team members is performed dynamically at run time. Upon role assignment, each robot i is assigned its tactic T_i from the current step of the sequence for that role.

A tactic, T , encapsulates a single robot behavior. Each robot i executes its own tactic as created by the current play P . A tactic T_i determines the skill state machine SSM_i to be executed by the robot i .

A skill, S , is a focused control policy for performing some complex action. Each skill is a member of one, or more, skill state machines SSM_1, SSM_2, \dots . Each skill S determines what skill it transitions to S' based upon the world state, the time skill S has been executing for, and the executing tactic for that robot.

We construct the robot cognition using a similar architecture. Plays, tactics, and skills, form a hierarchy for team control. Plays control the team behavior through tactics, while tactics encapsulate individual robot behavior and instantiate actions through sequences of skills. Skills implement the focused control policy for actually generating useful actions. Figure 1 shows the SSM s and transitions for an example tactic: *CatchKickToGoal*, which contains six skills. Each node in the figure is a skill and the edges show the transition between skills.

Segway soccer is a team sport, and therefore the building of our game strategy required not only execution of single robot behavior, but also coordination with the team member, the human player. The current coordination is simple and basically based upon two fixed plays for offensive and defensive situation respectively. Our offensive play is shown as follows, in which the termination condition is either play aborted or the situation changed (a turn-over of ball possession announced by the referee). There are two roles in this play, one passes the ball to the other who positions down field and waits for receiving a pass. Once the positioning




Figure 1: Skill state machines (SSMs) for an example tactic: *CatchKickToGoal*.

is done (e.g the other is closer to the opponent goal), the passing is performed.

```

PLAY Naive Offense
APPLICABLE offense
DONE aborted OR !offense
ROLE 1
  pass 2
  none
ROLE 2
  position_down_field
  receive_pass
  none
  
```

Our current coordination is purely observation based. Each player assigns role from his own eyeshot without communication. For example, should the robot think the team member is closer to the ball, the robot (ROLE 2) would choose to position and receive the ball from its team member (ROLE 1). Furthermore, the robot knows which side gains possession of the ball from the referee announcement, therefore it tells offensive from defensive situation clearly and thus it has deterministic idea of which play the team is using. The robot makes an assumption that its team member is performing the same game play as itself. The robot infers what tactic the team member is executing from the team play. For example, after receiving the ball from the team member, as a passer the robot would assume the team member go forward to a tactically advantageous position to receive a pass. The predefined play for team coordination helps motion modeling, which will be further discussed in section 3.

3. PLAY-BASED MOTION MODELING

In this section, we take a multi-target tracking problem as a detailed example to show the extension of the tactic-based motion modeling method in general when the team coordination knowledge (play) is incorporated. First we give an introduction of the environment and targets under the Segway soccer setup. Second, we describe detailed motion models for both the ball and the team member. Third, we extend the tactic-based motion modeling to the play level when both the ball and the team member are included into the tracking. We show how we model the play-dependent interactions between the team member, the robot and the ball and set up a base for giving the multi-model tracking algorithm in the next section. Although we present the Segway soccer domain as an example, the formalism is general.

3.1 Tracking Scenario

Many tracking scenarios involve multiple moving targets. In a Segway soccer game, we need to track the ball, the hu-

man team member and the two opponents. Each team is identified by their distinct color. The ball is orange [13]. An observation from the sensors might consist of multiple measurements due to the moving objects and the clutter. Generally the clutter has similar color as the targets we are interested in and it causes multiple hypothesis for the true targets. Therefore, we construct two multi-target trackers in the system, for the ball and the team member respectively. We use two separate trackers instead of one because we can differentiate the ball with the team member thanks to the color-based vision system. We use multi-target tracker instead of single target tracker because we can keep track of the true target and the false positive at the same time without losing any of them and perform ball(or team member) recognition from a pool of tracked objects later.

3.2 State Space and Dynamics

Let $\mathbf{x}_t = [\mathbf{x}_{1,t}^T, \mathbf{x}_{2,t}^T, \dots, \mathbf{x}_{K_t,t}^T]^T$, denote a combined target state vector for K_t unknown and time-varying number of targets. The general parameterized system process for the k th target $\mathbf{x}_{k,t}$ at time t is given by:

$$\mathbf{x}_{k,t} = \mathbf{f}_k(\mathbf{x}_{k,t-1}, \mathbf{u}_{k,t-1}, \mathbf{v}_{k,t-1}) \quad (1)$$

where \mathbf{f}_k is the parameterized state transition functions for the k th target; \mathbf{x}, \mathbf{u} are the state and input vectors; \mathbf{v} is the process noise vectors of known statistics. Given the number of targets at $t-1$ and t , the state transition can be represented as follows:

$$p(\mathbf{x}_t | \mathbf{x}_{t-1}, K_t, K_{t-1}) = \begin{cases} p_0(\mathbf{x}_{K_t,t}) \prod_{k=1}^{K_{t-1}} p(\mathbf{x}_{k,t} | \mathbf{x}_{k,t-1}), & \text{if } K_t = K_{t-1} + 1 \\ \prod_{k=1}^{K_t} p(\mathbf{x}_{k,t} | \mathbf{x}_{k,t-1}), & \text{if } K_t = K_{t-1} \\ \prod_{k=1, k \neq k^*}^{K_{t-1}} p(\mathbf{x}_{k,t} | \mathbf{x}_{k,t-1}), & \text{if } K_t = K_{t-1} - 1 \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

The prior $p(\mathbf{x}_{k,t} | \mathbf{x}_{k,t-1})$ can be evaluated from (1). The function $p_0(\mathbf{x}_{K_t,t})$ is the initial distribution of a new target state and k^* is a target that vanishes at t . We assume at most one target appears or vanishes in each time step. The number of targets is obtained deterministically (see section 4.1 for the details).

3.3 Observation Model

We assume that measurements are independent of each other and either originate from true targets or clutter. Furthermore, each target generates at most one measurement in each time step but may not be detected by the sensor.

Let $\mathbf{z}_t = [\mathbf{z}_{1,t}^T, \mathbf{z}_{2,t}^T, \dots, \mathbf{z}_{M_t,t}^T]$ denote an observation vector for M_t measurements. The observation equation for modeling the c th measurement originating from the k th target is given as:

$$\mathbf{z}_{c,t} = \mathbf{h}_k(\mathbf{x}_{k,t}, \mathbf{n}_{c,t}) \quad (3)$$

where $\mathbf{n}_{c,t}$ is a zero-mean observation noise with a known covariance Σ_n . The measurement originating from a clutter is modelled to be uniformly distributed within the entire visible region (observation volume V) of the camera. We use a deterministic method to perform measurement-to-target association. Let α_t denote the association vector which indicate the measurement-to-target assignment, whose i th element $\alpha_{i,t}$ is set to j if the i th measurement originates from the j th target, or zero if it originates from clutter. Let N_{C_t} denote the number of clutter points, and Ω_D denote the

set of measurements indices corresponding to the detected targets, the likelihood function for measurement \mathbf{z}_t can be written as:

$$p(\mathbf{z}_t | \mathbf{x}_t, K_t, \alpha_t) = \left(\frac{1}{V}\right)^{N_{C_t}} \prod_{l \in \Omega_D} p(\mathbf{z}_{l,t} | \mathbf{x}_{\alpha_{l,t},t}) \quad (4)$$

where $p(\mathbf{z}_{l,t} | \mathbf{x}_{\alpha_{l,t},t})$ is evaluated from (3).

3.4 Ball Motion Modeling

In our Segway RMP soccer robot environment, we define five models to model the ball motion (for the rest of this paper, for simplicity, we use \mathbf{x}_t to represent the ball state, and use \mathbf{x}'_t to represent the team member state).

- *Free-Ball, Robot-Grab-Ball, Robot-Kick-Ball.* We use the same models as *Free-Ball, Grab-Ball* and *Kick-Ball* introduced in [7] respectively.
- *Human-Grab-Ball.* The ball is held by the team member. we can infer the ball position similarly if we know the team member position well.
- *Human-Kick-Ball.* The ball is kicked by the team member and it is supposed to be either a pass to the robot or a shoot at the goal .

In general, we will have n motion models m_1, m_2, \dots, m_n .

3.5 Team Member Motion Modeling

We define four models to model the human team member's motion.

- *Random Walk.* The team member is wondering in the field. So the state at the new time is the state at the current time with some additive zero-mean (assumed Gaussian) noise.
- *Holding Ball.* The team member is holding the ball without moving and waiting for the robot to receive the ball. Should the robot know the ball position well, it can infer the team member position by the ball position in a similar way as *Robot-Grab-Ball* for ball motion modeling.
- *Accelerating (Decelerating).* The team member dashes (stops) to obtain (lose) a velocity in a short time.
- *Positioning.* The team member is going to a predefined tactical position with a constant speed. This case happens mostly after the team member passing the ball to the robot and moving down the field toward opponent's goal.

3.6 Play Based Model Transitions

Given the knowledge of the team coordination plan (the play P_{t-1} at time $t-1$), the robot can infer what tactic the team member is executing (T'_{t-1}), which provides valuable information about the motion model of the team member (m'_t). Both the robot and the team member act over the ball in a Segway soccer game. The motion model of the ball (m_t) is therefore affected by what tactic the robot (T_{t-1}) and the team member (T'_{t-1}) are executing.

From the previous subsection, we know that the model index m determines the present model being used. For our team member tracking example, $m'_t = i, i = 1, \dots, 4$. In our




Figure 2: Play-Based motion modeling, where m_1, m_2, \dots, m_n are n models, P_a is the team play, v_b is the additional information. $h_{i,j}$ is the transition probability from model m_i to model m_j given m_i , and $\langle P_a, v_b \rangle$. Each layer in the graph is conditioned on a particular combination of the play executed and the additional information obtained.

approach, it is assumed that the team member motion model index, m'_t , conditioned on the previous tactic executed T'_{t-1} by the team member, and other useful information v'_t (such as ball state), is governed by an underlying Markov process, such that, the conditioning parameter can branch at the next time-step with probability.

$$p(m'_t = i | m'_{t-1} = j, T'_{t-1}, v'_t) = h'_{i,j} \quad (5)$$

where $i, j = 1, \dots, N_{m'}$. Since T'_{t-1} can be determined by P_{t-1} , we get

$$h'_{i,j} = p(m'_t = i | m'_{t-1} = j, P_{t-1}, v'_t) \quad (6)$$

Since we can draw $p(m'_t = i | m'_{t-1} = j)$ in an $N_{m'} \times N_{m'}$ table, we can create a table for Equation 6 with a third axis which is defined by the tuple $\langle P_a, v_b \rangle$ as shown in Figure 2. Here the play P_a , is the primary factor that determines whether m_i transits to m_j and what the transition probability is, while the information v_b determines the prior condition of the transition. Each layer in the graph is conditioned on a particular combination of the tactic executed and the additional information obtained.

For our ball tracking example, $m_t = i, i = 1, \dots, 5$. Similarly,

$$h_{i,j} = p(m_t = i | m_{t-1} = j, T_{t-1}, v_t) \quad (7)$$

where $i, j = 1, \dots, N_m$. Since T_{t-1}, T'_{t-1} can be determined by P_{t-1} , we get

$$h_{i,j} = p(m_t = i | m_{t-1} = j, P_{t-1}, v_t) \quad (8)$$

Suppose the current team play is the *Naive Offense* in Section 2.2, we can obtain the corresponding motion model transitions for both the ball and the team member using the play-based method (Figure 3).

4. MULTI-MODEL MOTION TRACKING

In this section, we first describe the clustering algorithm we used to continuously monitor the appearance and disappearance of regions of interest (ROIs) on the field. We then

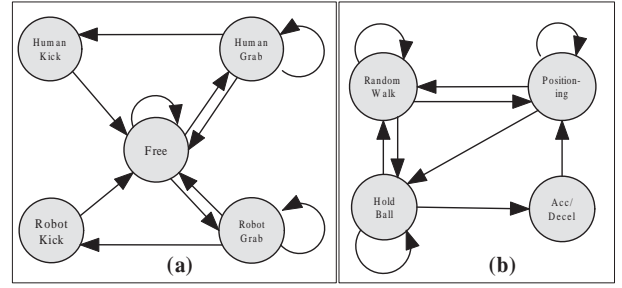


Figure 3: Object motion modeling based on the play: Naive Offense. Each node is a model. Models transit to one another according to the predefined probabilities (not shown in the figure). (a) Ball motion model. (b) Human team member motion model.

use dynamic Bayesian networks to represent the whole system. We give the detailed sequential Monte Carlo methods of tracking in a multi-model multi-hypothesis scheme finally.

4.1 Clustering Algorithm

The idea of this algorithm is to group a set of ROIs, $\mathcal{S}_t = \{S_t^{(j)}\}_{j=1}^{J_t}$, within a buffer of observations $\mathcal{Z}_t = \{\mathbf{z}_{t'}\}_{t'-\tau}^t$, where τ is the length of the sliding window [10]. If the current targets are widely separated, the measurements originating from them will be clustered in the locations where the targets had visited from $t - \tau$ to t . Therefore the distance between two measurements within two successive time steps can be taken as the clustering criteria. Let $d_{c,l}(t' + 1, t')$ denote the normalized distance between the c th and l th measurements of $\mathbf{z}_{t'+1}$ and $\mathbf{z}_{t'}$.

$$d_{c,l}(t' + 1, t') = e^{T_{c,l}(t' + 1, t')} \Sigma_e^{-1} e_{c,l}(t' + 1, t') \quad (9)$$

where $c \in \{1, \dots, M_{t'+1}\}$ and $l \in \{1, \dots, M_{t'}\}$, $e_{c,l}(t' + 1, t') = \mathbf{z}_{c,t'+1} - \mathbf{z}_{l,t'}$ and $\Sigma_e = 2\Sigma_n$. For all measurements of $\mathbf{z}_{t'+1}$, we have a set of normalized distances $\{d_{c,l}(t' + 1, t')\}_{c=1}^{M_{t'+1}}$. The c^* th measurement $\mathbf{z}_{c^*,t'+1}$ will be grouped with $\mathbf{z}_{l,t'}$ in the same cluster $S_t^{(j)}$ if

$$d_{c^*,l}(t' + 1, t') = \operatorname{argmin}\{d_{c,l}(t' + 1, t')\}_{c=1}^{M_{t'+1}} \quad (10)$$

and

$$d_{c^*,l}(t' + 1, t') \leq \epsilon_z \quad (11)$$

where ϵ_z is a threshold. If none of the measurements of $\mathbf{z}_{t'+1}$ can be grouped with $\mathbf{z}_{l,t'}$, the search procedure continues and uses $\mathbf{z}_{t'+b}, 1 < b \leq \tau$ until (10) and (11) are both satisfied. Meanwhile the algorithm propagates through time to group other measurements in the remaining observations. Since the detected regions obtained via (10) and (11) possibly arise from clutter points, another threshold is set to examine the number of measurements in the region to eliminate the false positives. That is, if the number of measurements in the j th region $S_t^{(j)}$ is less than τ_{min} , defined as the minimum number of clustered measurements in a region required to identify a target, it is discarded; otherwise, the j th region is classified as originating from a target. Once a set of ROIs $\mathcal{S}_t = \{S_t^{(j)}\}_{j=1}^{J_t}$ is obtained, it is necessary to determine which region belongs to which existing active track or a

new track deterministically. Let $\gamma_t = [\gamma_{1,t}, \gamma_{2,t}, \dots, \gamma_{K_t,t}]^T$ denote the track-to-region association vector. $\gamma_{k,t}$ is used to indicate the association between the ROIs and the active tracks. $\gamma_{k,t}$ is evaluated as j if track k can be associated with $S_t^{(j)}$, otherwise it is zero. Refer to [10] for more details of the clustering algorithm.

4.2 DBN Representation

Following the play-based motion model, we can use dynamic Bayesian networks (DBNs) to represent the whole system for team member and ball tracking in a natural and compact way as shown in Figure 4 and Figure 5 respectively. In the two graphs, the system state is represented by variables (play P , tactic T , infrared sensor measurement s , ball state \mathbf{x} , ball motion model index m , vision sensor measurement of ball \mathbf{z} , team member state \mathbf{x}' , team member motion model index m' , vision sensor measurement of team member \mathbf{z}'), where each variable takes on values in some space. The variables change over time in discrete intervals, so that e.g. \mathbf{x}_t is the ball state at time t .

Furthermore, the edges indicate dependencies between the variables. For instance, in Figure 5 the ball motion model index m_t depends on $m_{t-1}, T_{t-1}, T'_{t-1}, s_t$ and \mathbf{x}_{t-1} , hence there are edges coming from the latter five variables to m_t . For the rest of this section, we give the ball-tracking algorithm following Figure 5. The team-member-tracking algorithm can be obtained similarly following Figure 4.

4.3 Importance Sampling Function

We use the sequential Monte Carlo method to track the motion model m and the multi-target state \mathbf{x} . Particle filtering is a general purpose Monte Carlo scheme for tracking in a dynamic system [6]. It maintains the belief state at time t as a set of particles $p_t^{(1)}, p_t^{(2)}, \dots, p_t^{(N_s)}$, where each $p_t^{(i)}$ is a full instantiation of the tracked variables $\{p_t^{(i)}, w_t^{(i)}\}$, $w_t^{(i)}$ is the weight of particle $p_t^{(i)}$ and N_s is the number of particles. In our case, $p_t^{(i)} = \langle \mathbf{x}_t^{(i)}, m_t^{(i)} \rangle$. To make our notation more concrete, a particular particle $p_t^{(i)}$, which is tracking K_t multi-target state vector \mathbf{x}_t and motion model m_t , is given as [8]:

$$p_t^{(i)} = \begin{pmatrix} m_{1,t}^{(i)} & m_{2,t}^{(i)} & \dots & m_{K_t,t}^{(i)} \\ x_{1,t}^{(i)} & x_{2,t}^{(i)} & \dots & x_{K_t,t}^{(i)} \\ y_{1,t}^{(i)} & y_{2,t}^{(i)} & \dots & y_{K_t,t}^{(i)} \\ \dot{x}_{1,t}^{(i)} & \dot{x}_{2,t}^{(i)} & \dots & \dot{x}_{K_t,t}^{(i)} \\ \dot{y}_{1,t}^{(i)} & \dot{y}_{2,t}^{(i)} & \dots & \dot{y}_{K_t,t}^{(i)} \end{pmatrix} \quad (12)$$

We sample the ball motion model following the ball-tracking DBN as below:

$$m_{k,t}^{(i)} \sim p(m_{k,t} | m_{k,t-1}^{(i)}, \mathbf{x}_{k,t-1}^{(i)}, s_t, T_{t-1}, T'_{t-1}) \quad (13)$$

Note that T_{t-1} and T'_{t-1} are inferred deterministically from P_{t-1} instead of sampling. Conditioned on the ball motion model $m_{k,t}^{(i)}$, we then use the importance function introduced in [10] to sample ball state $\mathbf{x}_{k,t}^{(i)}$:

$$\mathbf{x}_{k,t}^{(i)} \sim \begin{cases} q_D(\mathbf{x}_{k,t} | m_{k,t}^{(i)}, \mathbf{x}_{k,t-1}^{(i)}), & k \notin \mathcal{K}_{S_t} \\ q_{DS}(\mathbf{x}_{k,t} | m_{k,t}^{(i)}, \mathbf{x}_{k,t-1}^{(i)}, S_t^{(j)}), & k \in \mathcal{K}_{S_t} \end{cases} \quad (14)$$

where $k \in \mathcal{K}_{S_t}$ are those tracks with $\gamma_{k,t} = j$ and $j \in \Omega_D$, and $q_D(\cdot)$ and $q_{DS}(\cdot)$ are the proposal functions for $x_{k,t}$

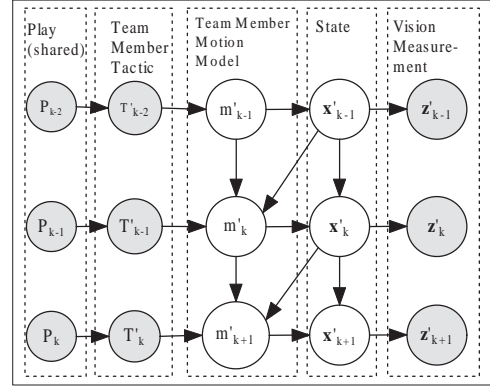


Figure 4: A dynamic Bayesian network for team member tracking with a Segway RMP robot. Filled circles represent deterministic variables which are observable or are known as the tactic or the play that the robot is executing.

without and with an associated ROI $S_t^{(j)}$, given as follows, respectively,

$$q_D(\mathbf{x}_{k,t} | m_{k,t}^{(i)}, \mathbf{x}_{k,t-1}^{(i)}) = p(\mathbf{x}_{k,t} | m_{k,t}^{(i)}, \mathbf{x}_{k,t-1}^{(i)}) \quad (15)$$

$$q_{DS}(\mathbf{x}_{k,t} | m_{k,t}^{(i)}, \mathbf{x}_{k,t-1}^{(i)}, S_t^{(j)}) =$$

$$\mu p(\mathbf{x}_{k,t} | m_{k,t}^{(i)}, \mathbf{x}_{k,t-1}^{(i)}) + (1 - \mu) q(\mathbf{x}_{k,t} | m_{k,t}^{(i)}, \mathbf{x}_{k,t-1}^{(i)}, S_t^{(j)}) \quad (16)$$

where $0 \leq \mu \leq 1$ and $q(\cdot)$ is a uniform sampling from the associated ROI $S_t^{(j)}$. If $\mu = 1$, the importance sampling function is reduced back to the dynamic prior. If $\mu = 0$, all particles are generated from the data-dependent importance function. If $0 < \mu < 1$, this proposal combines the dynamic prior and the current ROIs to generate representative particles.

4.4 Birth, Death and Update Moves

Assuming that there are K_t^b ROIs that cannot be associated with any existing track, we will initiate a new track in each time step from one of these regions instead of initiating K_t^b tracks simultaneously in order to fit the birth move with the assumed system process model in (2). When an existing track cannot be associated with a region at a given time, the target being tracked by the tracker may have disappeared or temporarily experience a short period of data loss. Thus we may remove the track for the target only if it has failed to associate with any ROI with τ_d time steps. Refer to [10] for the detailed algorithm of birth move and death move.

In the update move, there is no change in terms of the number of ROIs. We only need to update the target states with a common value of number of targets $K_t^{(i)} = K_{t-1}^{(i)}$, using the sequential importance sampling method as follows:

$$SIS[\{\mathbf{x}_{t-1}^{(i)}, m_{t-1}^{(i)}, w_{t-1}^{(i)}\}_{i=1}^{N_s}, \mathbf{z}_t, s_t, T_{t-1}, T'_{t-1}] =$$

- 01 for $i = 1 : N_s$
- 02 for $k = 1 : K_t$
- 03 draw $m_{k,t}^{(i)} \sim p(m_{k,t} | m_{k,t-1}^{(i)}, \mathbf{x}_{k,t-1}^{(i)}, s_t, T_{t-1}, T'_{t-1})$.
- 04 if track k has corresponding ROI $S_t^{(j)}$
- 05 draw $\mathbf{x}_{k,t}^{(i)} \sim q_{DS}(\mathbf{x}_{k,t} | m_{k,t}^{(i)}, \mathbf{x}_{k,t-1}^{(i)}, S_t^{(j)})$

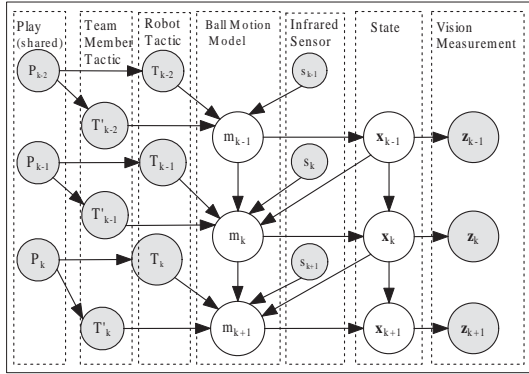


Figure 5: A dynamic Bayesian network for ball tracking with a Segway RMP robot.

```

06   else
07     draw  $\mathbf{x}_{k,t}^{(i)} \sim q_D(\mathbf{x}_{k,t} | m_{k,t}^{(i)}, \mathbf{x}_{k,t-1}^{(i)})$ 
08   end if
09 end for
10 set  $w_t^{(i)} = w_{t-1}^{(i)} \cdot p(\mathbf{z}_t | x_t, K_t, \alpha_t)$ 
11 end for
12 Calculate total weight:  $w = \sum_{i=1}^{N_s} [w_t^{(i)}]$ 
13 for  $i = 1 : N_s$ 
14   Normalize:  $w_t^i = w_t^i / w$ 
15 end for
16 Resample.

```

The inputs of the algorithm are samples drawn from the previous posterior $\langle \mathbf{x}_{t-1}^{(i)}, m_{t-1}^{(i)}, w_{t-1}^{(i)} \rangle$, the present vision and infrared sensory measurement \mathbf{z}_t, s_t , the robot's tactic T_{t-1} , and the team member's tactic T'_{t-1} . The outputs are the updated weighted samples $\langle \mathbf{x}_t^{(i)}, m_t^{(i)}, w_t^{(i)} \rangle$. In the sampling algorithm, first, a new ball motion model index, $m_t^{(i)}$, is sampled according to (13) at line 03. Then given the model index, and previous ball state, a new ball state is sampled according to (14) at line 05/07. According to (4), the importance weight of each sample is given by the likelihood of the vision measurement given the predicted new ball state at line 10. Finally, each weight is normalized and the samples are resampled. Then we can estimate the ball state based on the mean of all the $\mathbf{x}_t^{(i)}$. Though we are trying to eliminate the clutter from the beginning of tracking (clustering algorithm), due to the property of the multi-target tracker, further recognition process might be done in order to figure out which tracked target is the true ball. Similarly the state of the team member \mathbf{x}'_t can be obtained from the team member tracker.

5. EXPERIMENTAL RESULTS

From previous work we knew the initial speed and accuracy of the ball velocity after a kick motion. We profiled the system and measurement noise as well. In this section, we evaluate the effectiveness of our tracking system in both simulated and real-world tests.

5.1 Simulation Experiments

Because it is difficult to know the ground truth of the target's position and velocity in the real robot test, we do the simulation experiments to evaluate the precision of tracking.

Table 1: The average RMS error of position estimation and velocity estimation from human trackers and ball trackers.

Motion Model	Single Model	Multiple Model
Human Position Est RMS (m)	0.0030	0.0014
Human Velocity Est RMS (m/s)	0.42	0.025
Ball Position Est RMS (m)	0.0028	0.0017
Ball Velocity Est RMS (m/s)	0.4218	0.0597

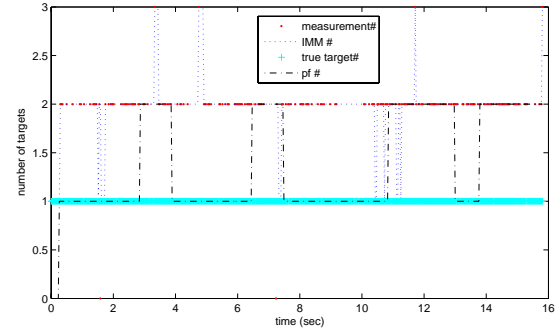


Figure 6: A comparison of the target detection performance between the proposed method and the IMM tracker when only one target exists with surrounding clutters.

Experiments are done following the *Naive Offense* play, in which the robot acts as the receiver and the human team member acts as the passer. Noises are simulated according to the model we profiled in previous work. In the beginning, the team member holds the ball. After a fixed amount of time, the ball is kicked towards the robot, and the team member moves forward to a predefined location.

We implement both a single model tracker and a play-based multi-model tracker for the ball and the team member. We simulate the experiment for 50 runs, and then compare the performance of the two trackers with different implementations. The average RMS error of position estimation and velocity estimation are shown in Table 1. The results show that the play-based multi-model scheme performs much better than the single model especially in terms of velocity estimation. Because with the play-based motion model, when the ball is being kicked, most particles evolving using the transition model determined by the play will change its motion model $m_t^{(i)}$ from *Free-Ball* to *Human-Kick-Ball*, and a velocity will be added to the ball accordingly.

5.2 Multi-Target Tracking Test

In this test, one Segway RMP robot is tracking one or more balls on the field with *SearchBall* tactic. We would like to compare solely the target detection performance between the proposed method and the IMM tracker. A scenario with K_t ($0 \leq K_t \leq 3$) balls appearing and disappearing at different times and there are a set of false positives at fixed position in the surroundings.

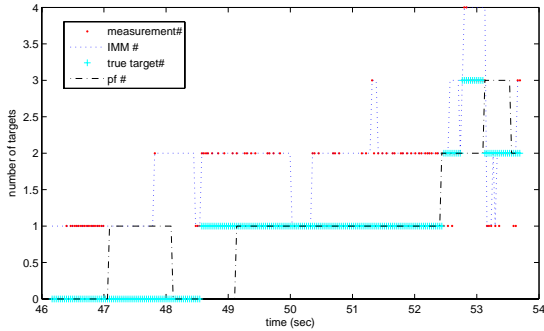


Figure 7: A comparison of the target detection performance between the proposed method and the IMM tracker when multiple targets exist with surrounding clutters.

When estimating the number of targets, 3600 particles are used in the proposed method. Figures 6-7 summarize the results. In both figures, the dots show the number of measurements at a given time. The dotted line represents the number of the targets tracked by the IMM tracker. The dashdotted line represents the number of targets tracked by the multi-model multi-target tracker proposed in this paper. The crosses show the true number of the targets at any given time. As shown in the figure, the IMM tracker is sensitive to the number of measurements, while our approach is more robust and consistent to high clutter density. Since the detection is basically performed on the clustering of the observations and the association between the detected ROIs and the existing tracks, it is computational low-cost. Therefore it is also practical for real-time multi-target detection and tracking.

5.3 Team Cooperation Test

We do experiments on the Segway RMP soccer robot executing the offensive play and coordinating with the human team member. The test setup is demonstrated in Figure 8, in which the digits along the lines show the sequence of the whole strategy, the filled circle at position B represents the robot, the unfilled circle at position E represent an opponent player, and the shaded circle represent the human team member.

When each run begins, the human team member is at position A . With this team cooperation plan (play), the robot chooses the tactic *CatchKickToTeammember* to execute, in which the robot starts with the skill *Search-Ball*. When the robot finds the ball, the team member passes the ball directly to the robot and chooses a positioning point to go to either at C or D . The robot grabs the ball after the ball is in the catchable area and is detected by the infrared sensor (skill *Grab-Ball*). Next the robot searches for the team member holding the ball with its catcher (skill *Search-Teammember*). After the robot finds the team member, the robot kicks the ball to its team member (skill *KickToTeammember*) and the team member shoots at the goal, completing the whole offensive play. Each run ends in one of the following conditions.

- succeed if the human receives the ball from the robot or the human does not receiver the ball but the pass




Figure 8: A demonstration of a naive team cooperation plan in offensive scenario. The digits along the lines show the sequence of the whole plan. The filled circle at position B represents the robot. The unfilled circle at position E represent an opponent player. The shaded circle represent the human team member.

Table 2: The average time taken over all the successful runs.

Motion Model	Single Model	Multi-Model
Mean Time (sec)	33.4	22.6

can be considered as a “good” one.

- fail if the robot is in searching for the ball or the team member for more than 30 seconds.
- fail if the ball is outside the field before the robot catches it.

In the experiment over 15 runs, the robot with single model trackers fails 5 of the total. While the robot with play-based multi-model trackers fails 2 of the total. We also keep track of the mean time taken in all the successful runs. We list the result in Table 2. Using play-based multi-model tracking saves 32.3% time in terms of completing the whole play over single model tracking. During the experiment, we note that when using the single model tracking, most time were spent on searching the team member. Incorporating the team cooperation knowledge known as play into the team member motion modeling greatly improves the accuracy of the team member motion model and therefore avoids taking time in searching a lost target from scratch.

6. RELATED WORK

Tracking moving targets using a Kalman filter is the optional solution if the system follows a single model, \mathbf{f} and \mathbf{h} in Equation 1 and 3 are known linear functions and the noise \mathbf{v} and \mathbf{n} are Gaussians [1]. Multiple model Kalman filters such as Interacting Multiple Model (IMM) are known to be superior to the single Kalman filter when the tracked target is maneuvering [3]. For nonlinear systems or systems with non-Gaussian noises, a further approximation is introduced, but the posterior densities are therefore only locally accurate and do not reflect the actual system densities.

Since the particle filter is not restricted to Gaussian densities, a tactic-based motion modeling method is proposed

in [7]. Based on that approach, we further introduce the play-based motion modeling method when team coordination knowledge is available. Another related approach was proposed to track a moving target using Rao-Blackwellised particle filter [9], in which a fixed transition table was used between different models. Our transition model is dependent on the play that the robot is executing and the additional information that matters. This model can be flexibly integrated into our existing STP architecture.

There have been different strategies in multi-target tracking. In order to handle the data association and tracking problem, the classical Joint Probabilistic Data Association Filter (JPDAF) adopts the methods like the extended Kalman Filter (EKF) for multi-target state estimation, whose tracking performance is known to be limited by the linearity of the data models [2]. Another approach known as sequential Monte Carlo methods are able to perform well even when the data models are nonlinear and non-Gaussian. However, almost all of these methods assume that the knowledge of true targets (without clutter) is given, which is not applicable in the field that Segway RMP soccer robots operates in.

Recently, a hybrid approach for online joint detection and tracking for multiple targets was proposed [10]. This approach does not rely on the clutter-free assumption. In this paper, based on their approach, we present a play-based multi-target tracking algorithm, which incorporates tactic information to eliminate the false alarms and to improve resampling efficiency. Compared to our method, first, existing techniques consider less complex dynamic systems where only one part of the state space is non-linear. In contrast, our approach estimate a system where multiple components are highly non-linear (Segway RMP robot motion, ball motion, team member motion). Second, most existing techniques examine their performance with simulated experiments, while we test our approach in real robot experiments. Third, our approach goes beyond existing techniques by incorporating team cooperation information into the tracking process which further improves the performance.

7. CONCLUSIONS AND FUTURE WORK

Motivated by the interactions between a team and the tracked target, we contribute a method to achieve efficient tracking through using a play-based motion model and combined vision and infrared sensory information. This method gives the robot a more exact task-specific motion model when executing different tactics over the tracked target (e.g. the ball) or collaborating with the tracked target (e.g. the team member). Then we represent the system in a compact dynamic Bayesian network and use particle filter to keep track of the motion model and target state through sampling. The empirical results from the simulated and using the real robot agent show the efficiency of the multi-model tracking over single model tracking.

If the teammate is a human, not a robot, the certainty that the teammate is executing the expected play or tactic could be reduced. That is, the human teammate could fail to execute the desired play or tactic. Future work will take such uncertainty into account. A better human team member modeling (for example, include intercepting the moving ball, mark a player, covering the goal) will also help. Another interesting work is to know how the performance of the presented method is affected by the presence of tactics

of the team member that are not exactly determined in the team coordination plan.

8. ACKNOWLEDGMENTS

We would like to thank the members of the CMBalance Segway soccer team for their help with developing the infrastructure for the Segway robots. In particular, we thank Brett Browning for the vision and initial control systems and Brenna Argall for the game skills.

This work was supported by United States Department of the Interior under Grant No. NBCH-1040007. The content of the information in this publication does not necessarily reflect the position or policy of the Defense Advanced Research Projects Agency (DARPA), US Department of Interior, US Government, and no official endorsement should be inferred.

9. REFERENCES

- [1] S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A tutorial on particle filters for on-line non-linear/non-gaussian bayesian tracking. *IEEE Transactions on Signal Processing*, 50(2):174–188, Feb. 2002.
- [2] Y. Bar-Shalom and T. E. Fortmann. *Tracking and Data Association*. Academic Press, Inc, 1988.
- [3] Y. Bar-Shalom, X.-R. Li, and T. Kirubarajan. *Estimation with Applications to Tracking and Navigation*. John Wiley & Sons, Inc, 2001.
- [4] B. Browning, J. Bruce, M. Bowling, and M. Veloso. STP: Skills, tactics and plays for multi-robot control in adversarial environments. *IEEE Journal of Control and Systems Engineering*, 219:33–52, 2005.
- [5] B. Browning, J. Searock, P. E. Rybski, and M. Veloso. Turning segways into soccer robots. *Industrial Robot*, 32(2):149–156, 2005.
- [6] A. Doucet, N. D. Freitas, and N. Gordon, editors. *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York, 2001.
- [7] Y. Gu. Tactic-based motion modelling and multi-sensor tracking. *Proceedings of Twentieth National Conference on Artificial Intelligence*, 2005.
- [8] C. Kreucher, K. Kastella, and A. O. H. III. Multi-target sensor management using alpha-divergence measures. pages 209–222, 2003.
- [9] C. Kwok and D. Fox. Map-based multiple model tracking of a moving object. *Proceedings of eight RoboCup International Symposium*, July 2004.
- [10] W. Ng, J. Li, S. Godsill, and J. Vermaak. A hybrid approach for online joint detection and tracking for multiple targets. *IEEE Aerospace Conferences*, 2005.
- [11] D. Schulz, W. Burgard, and D. Fox. People tracking with mobile robots using sample-based joint probabilistic data association filters. *International Journal of Robotics Research*, 22(2), 2003.
- [12] J. Searock, B. Browning, and M. Veloso. Turning Segways into Soccer Robots. In *Proceedings of IROS'04*, Sendai, Japan, September 2004.
- [13] M. Veloso, B. Browning, P. Rybski, and J. Searock. Segwayrmp robot football league rules. Technical report, <http://www.cs.cmu.edu/robosoccer/segway/>, 2005.