Decision-Theoretic Approach to Maximizing Fairness in Multi-Target Observation in Multi-Camera Surveillance

(Extended Abstract)

Prabhu Natarajan, Kian Hsiang Low, and Mohan Kankanhalli
Department of Computer Science, National University of Singapore, Republic of Singapore
{prabhu, lowkh, mohan}@comp.nus.edu.sg

ABSTRACT

Central to the problem of active multi-camera surveillance is the fundamental issue of fairness in the observation of multiple targets such that no target is left unobserved by the cameras for a long time. To address this important issue, we propose a novel principled decision-theoretic approach to control and coordinate multiple active cameras to achieve fairness in the observation of multiple moving targets.

Categories and Subject Descriptors

I.4.8 [**Scene Analysis**]: Tracking; I.2.9 [**Robotics**]: Commercial robots and applications, Sensors

General Terms

Algorithms, Performance, Experimentation, Security

Keywords

Surveillance and security; robot teams; multi-robot systems; robot coordination

1. INTRODUCTION

Active cameras are increasingly used in surveillance for monitoring and tracking targets in high-resolution images/videos. These cameras are commonly known as PTZ (pan-tilt-zoom) cameras and are needed to be controlled and coordinated efficiently to achieve a desired surveillance task. Recent works in controlling and coordinating active cameras are designed to maximize the number of targets to be observed in active cameras [5, 6] and to observe certain targets at a desired resolution [2, 7]. When the cameras are controlled to achieve these surveillance tasks, one or more targets may receive less attention by the active cameras. In the worst case, these targets may not be observed at all. These targets that are not observed by the active cameras for long duration may be adversarial and cause potential threats to the environment. Therefore, it is necessary to achieve fairness in observation of targets in active camera surveillance. That is, we need to maximize the observations of those targets that are not observed or least observed by the active cameras. In this paper, we propose a novel decision-theoretic approach to control and coordinate multiple active cameras to maximize the fairness in observation of multiple targets. Our decision-theoretic approach is based on Markov Decision Process (MDP) framework and max-min fairness metric.

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

2. MDP FRAMEWORK

Our surveillance setup consists of m targets, n active cameras, one or more static cameras, and a MDP controller. The targets are moving objects whose motion is stochastic in nature and active cameras are PTZ cameras that can obtain high-resolution images of the targets. The static cameras observe the surveillance environment at a low resolution. These static cameras are assumed to be calibrated and can obtain the 3D location, direction, and speed information of the targets to be passed to our MDP controller. Based on these information, our MDP controller computes the optimal PTZ actions of the active cameras to coordinate them in achieving our surveillance objective. The MDP controller is defined as a tuple $(\mathcal{S}, \mathcal{A}, T_f, R)$ that consists of the following:

States: Let $\mathcal{S}$ be set of joint states of targets and active cameras in the surveillance environment such that a joint state $S \triangleq (T_{\mathcal{M}}, C) \in \mathcal{S}$ consists of a pair of joint states $T_{\mathcal{M}} \in \mathcal{T}^m$ of $m = |\mathcal{M}|$ targets and $C \in \mathcal{C}^n$ of n active cameras. The sets $\mathcal{T}$ and $\mathcal{C}$ denote all possible states of each target and each active camera, respectively, and the set $\mathcal{M}$ denotes indices of the m targets. So, $\mathcal{S} = \mathcal{T}^m \times \mathcal{C}^n$. Let $T_{\mathcal{M}} \triangleq (t_1, t_2, \ldots, t_m) \in \mathcal{T}^m$ and $C \triangleq (c_1, c_2, \ldots, c_n) \in \mathcal{C}^n$ where $t_k \in \mathcal{T}$ and $c_i \in \mathcal{C}$ denote the states of target k and camera i, respectively. Let $t_k \triangleq (t_{l_k}, t_{d_k}, t_{v_k}, t_{o_k}) \in \mathcal{T}_l \times \mathcal{T}_d \times \mathcal{T}_v \times \mathcal{T}_o$ where $t_{l_k}, t_{d_k}, t_{v_k}$, and t_{o_k} denote target k's location, direction, speed, and observation time, respectively. So, $\mathcal{T} = \mathcal{T}_l \times \mathcal{T}_d \times \mathcal{T}_v \times \mathcal{T}_o$.

The state space $\mathcal{C}$ of an active camera is a finite set of discrete PTZ positions. Let $fov(C) \subset \mathcal{T}_l$ be a subset of target locations lying within the fov of all the cameras in state C.

Actions: The joint actions of the active cameras are PTZ commands that move the corresponding cameras to their specified states. Let a joint action of the n cameras be denoted by $A \triangleq (a_1, a_2, \ldots, a_n) \in \mathcal{A}$ where a_i denotes the PTZ command of camera i.

Transition Function: The transition function $T_f : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to [0, 1]$ models the probability $P(S'|S, A)$ of changing from the current joint state $S \in \mathcal{S}$ to the next joint state $S' \in \mathcal{S}$ using the joint action $A \in \mathcal{A}$. By exploiting the state transition dynamics of the surveillance environment, the transition model T_f can be factored into transition models of individual targets and active cameras [5, 6] as follows:

$$P(S'|S,A) = \begin{cases} \prod_{k=1}^{m} P(t'_k|t_k,C') & \text{if } P(c'_i|c_i,a_i)=1 \text{ for } i=1,\ldots,n, \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

where the transition model of individual target is given by

$$P(t'_k|t_k,C') = P(t'_{l_k}|t_{l_k},t'_{d_k},t'_{v_k})P(t'_{d_k}|t_{d_k})P(t'_{v_k}|t_{v_k}) \; P(t'_{o_k}|t'_{l_k},t_{o_k},C') \ .$$

The transition probabilities $P(t'_{l_k}|t_{l_k},t'_{d_k},t'_{v_k})$, $P(t'_{d_k}|t_{d_k})$, and $P(t'_{v_k}|t_{v_k})$ are Gaussian distributions, as reported in [5]. The transition probability for target's observation time is given by

$$P(t'_{o_k}=t_{o_k}+1|t'_{l_k},t_{o_k},C') = \begin{cases} 1 & \text{if } t'_{l_k} \in fov(C'), \\ 0 & \text{otherwise.} \end{cases}$$

$$P(t'_{o_k}=t_{o_k}|t'_{l_k},t_{o_k},C') = \begin{cases} 1 & \text{if } t'_{l_k} \notin fov(C'), \\ 0 & \text{otherwise.} \end{cases}$$

Objective/Reward Function: Let $R : \mathcal{S} \rightarrow \mathbb{R}$ be a real-valued reward function that represents the surveillance goal. Our goal is to maximize the fairness in observation of multiple targets in active cameras, which is defined in terms of the max-min fairness metric. Supposing the transition models of all targets are deterministic, we define a reward function R that measures the minimum observation time over all targets:

$$R(S) = R((T_{\mathcal{M}}, C)) \triangleq \min_{k\in\mathcal{M}} t_{o_k} \ . \tag{2}$$

Policy Computation: A policy function $\pi(S)$ in the MDP controller maps from each state S to a joint action A of cameras, that is, $\pi : S \rightarrow A$. Since the states of the targets in the next time step are uncertain due to stochasticity of their motion, we compute optimal policy $\pi^*(S)$ that maximizes the expected minimum observation time over all targets in the next time step:

$$A^* = \arg\max_{A\in\mathcal{A}} \sum_{T'_{\mathcal{M}}\in\mathcal{T}^m} R((T'_{\mathcal{M}}, C'))P(T'_{\mathcal{M}}|T_{\mathcal{M}}, C') \tag{3}$$

where $T'_{\mathcal{M}}$ and C' are, respectively, the joint states of the targets and active cameras in the next time step. Computing optimal policy using (3) incurs time that is exponential in number m of targets. We exploit the conditional independence property in the transition model to reduce the exponential computation time and, as a result, the policy function $\pi(S)$ (3) can be simplified to

$$A^* = \arg\max_{A\in\mathcal{A}} \prod_{k\in\mathcal{Y}} \sum_{t'_k\in\mathcal{T}_{C'}} P(t'_k|t_k, C') \ . \tag{4}$$

where $\mathcal{Y} \subseteq \mathcal{M}$ denotes the set of indices of all targets with minimum observation time in the current time step and $\mathcal{T}_{C'}$ denotes the set of a target's states whose locations are observed by the active cameras in their joint state C'. Policy computation time in (4) is linear in number m of targets.

Computing the policy in (4) only needs to consider all targets with minimum observation time. These targets may be beyond the fov's of some active cameras due to their spatial localities. To remedy this situation, the key idea is to repeatedly refine the set of optimal joint actions by preserving fairness in the observation of the remaining targets using (4) after ignoring those with minimum observation times.

3. RESULTS AND CONCLUSIONS

The proposed MDP Framework has been evaluated under various experimental setups [5, 6] like hall, corridor and junction. The results show that our approach maximizes the fairness in observation of targets efficiently in all three setups, as compared to the work in [5].

To summarize, we have proposed a decision-theoretic approach to control and coordinate multiple active cameras to improve the fairness in observation of targets. Our approach is based on MDP framework whose reward function is formulated as a max-min fairness metric. In future, we would like to improve fairness by exploiting the interactions of the targets [1] in the environment and also consider fairness in coordinating a team of mobile robots [3, 4] in surveillance.

4. ACKNOWLEDGMENTS

This research was carried out at the NUS-ZJU Sensor-Enhanced Social Media (SeSaMe) Center. It is supported by the Singapore National Research Foundation under its International Research Center @ Singapore Funding Initiative and administered by the Interactive Digital Media Programme Office.

5. REFERENCES

[1] T. Gan, Y. Wong, D. Zhang, and M. S. Kankanhalli. Temporal encoded F-formation system for social interaction detection. In *Proc. ACM Multimedia*, pages 937–946, 2013.

[2] N. Krahnstoever, T. Yu, S. N. Lim, K. Patwardhan, and P. Tu. Collaborative Real-Time Control of Active Cameras in Large Scale Surveillance Systems. In *Proc. M2SFA2*, pages 1–12, 2008.

[3] K. H. Low, W. K. Leow, and M. H. Ang Jr. Task allocation via self-organizing swarm coalitions in distributed mobile sensor network. In *AAAI*, pages 28–33, 2004.

[4] K. H. Low, W. K. Leow, and M. H. Ang Jr. Autonomic mobile sensor network with self-coordinated task allocation and execution. *IEEE Trans. Syst., Man, Cybern. C*, 36(3):315–327, 2006.

[5] P. Natarajan, T. N. Hoang, K. H. Low, and M. Kankanhalli. Decision-theoretic approach to maximizing observation of multiple targets in multi-camera surveillance. In *Proc. AAMAS*, pages 155–162, 2012.

[6] P. Natarajan, T. N. Hoang, K. H. Low, and M. Kankanhalli. Decision-theoretic coordination and control for active multi-camera surveillance in uncertain, partially observable environments. In *Proc. ICDSC*, pages 1–6, 2012.

[7] F. Z. Qureshi and D. Terzopoulos. Planning ahead for PTZ camera assignment and handoff. In *Proc. ICDSC*, pages 1–8, 2009.