

Cooperating Smart Cameras

Tiago Marques

Instituto Superior Técnico / UTL, Lisbon, Portugal

tiago.d.oliveira.marques@tecnico.ulisboa.pt

I. ABSTRACT

The increase need of surveillance in public places and recent technological advances on embedded video compression and communications made camera networks ubiquitous. There are however missing methodologies for watching so much data captured by so many cameras having few staff. In order to effectively track multiple targets, autonomous active camera networks require adequate detection, tracking, scheduling and control methodologies. Detection and tracking algorithms extract features from captured imagery which are used to build state posterior of the visual targets. Scheduling algorithms assign visual targets to cameras. Control methodologies set precise orientation and zoom references of the cameras

In this thesis, Information Theory elements are used in order to solve the scheduling and control problem. Each observable target in the environment corresponds to a source of information for which an observation corresponds to a reduction of the uncertainty and, as such, a gain in the information. Observation functions are shown to help avoiding extreme zoom levels while keeping smooth information gains.

At last, an agile optimization procedure proposal is made, in order to maximize the information gain for camera networks.

II. INTRODUCTION

The increased need for surveillance in public places and recent technological advances in embedded video compression and communications have made camera networks ubiquitous. However, at the moment, there are still missing suitable algorithms that are capable of automatically processing so much data captured by so many cameras having few staff members.

One of the issues associated with this problem is the decision on which control action to send to the pan-tilt-zoom cameras (PTZ) in order to successfully carry on the desired surveillance tasks. In simple words, the cooperative problem of controlling a network of PTZ cameras for the purpose of active surveillance in a dynamic scenario is that one wants to maintain high zoom levels without losing track on the targets in the scene.

The first autonomous surveillance systems were composed of multiple static cameras, working together to solve some practical tasks of interest such as tracking moving objects. The need for considering overlapping and wide fields (resulting in low-resolution images) of views of deployed cameras has led to the deployment of pan-tilt-zoom cameras in modern surveillance systems.

A number of new architectures appeared, such as master-slave camera configurations, and cooperative smart networks [11]. In the master-slave configuration, static cameras are used for event detection in order to direct the PTZ camera to the target of interest. In contrast, with more complex architecture, both static and PTZ camera streams are used for event analysis. This way, the global state of these systems is composed of both the individual state of each target and the camera state. In both architectures, there is a need for efficient and reliable target tracking methodologies and, in more complex cooperative architectures, there is also a need for camera management methodologies, in order to compute the optimal configuration of the network in order to coordinately maximize the coverage of tracked visual targets in the scene.

A. Related Work

The surveillance problem in active camera networks can be divided into two main components. The detection and tracking of visual targets of interest in the scene and the computation of a scheduling

policy to control each of the cameras' degree of freedom, in order to take into account the dynamics of the scene.

There are many generic methodologies used for target tracking. The commonly taken approach, based on the Bayes filter [1], [12], [11], is characterized at each iteration by the update of the state estimate based on the predicted state given by the motion model of the target and the observations given by the sensor. Another approach is taken in [8], [7], [3] and [2], where, contrary to the recursive approach, the trajectory is estimated in batches, making available both past and future observations for the estimation of the trajectory at a given time instance.

Regarding camera scheduling, Starzyk *et al.* [13] proposed a complete system for tracking multiple targets using cooperative cameras. The conflicts in behaviors are resolved using a central processor, which combines the individual desired behaviors in a single behavior, which reflects the best compromise between all of them.

The Multi-armed bandit algorithm is introduced in [14], [10] and [9] not as a camera management system, but a decision methodology for coordinately allocating *resources* to *projects*, e.g. robots that can travel to certain locations in order to discover events or network packets that can be routed to various channels in order to maximize the throughput. This framework can be used to model the problem of camera management, as well. Each camera is considered as a resource and each target as a project. The objective is to allocate cameras (resources) to targets (projects) in order to maximize some measure of reward over time.

Another approach, based on the information theory framework, is applied to the problem of camera management in [12] and [11]. This approach is based on the previous work on automatic zoom selection in [6], in which the camera parameters (i.e. the target-camera assignment) are chosen based on the mutual information gain between the state estimate and the estimate given an observation. In this approach, the Extended Kalman Filter is used as the selected tracking algorithm.

The aforementioned approach is the one adopted in this work. However, it is noteworthy that the design choice of the observation functions is still an open and ongoing challenge with the family of the information theory approaches. In the following

sections, we will present our particular approach.

B. Problem Formulation

A set of pan-tilt-zoom cameras is supposed to track and maintain trajectories of as many targets as possible. The targets are circulating in the environment.

We use an Information Theory framework in which the optimal control policy a^* for each camera is defined by the following maximization problem

$$a^* = \arg \max_a I(x; o) = \arg \max_a H_a(x) - H_a(x|o), \quad (1)$$

where $I(x; o)$ denotes the mutual information gain between the state estimate and the observation, and $H_a(x|o)$ is the conditional entropy of the state estimate given the observation, given that some control a was sent to the camera.

Each target has an assigned Extended Kalman Filter, with the motion model dependent on the targets being tracked and the pinhole camera model as sensor model. When a target is in the field of view of a camera, a new observation is available and its value is used to update the EKF. In this scenario, there is a reduction of the uncertainty in the target's state and, as such, a positive information gain. On the other hand, when a target is not in the camera's field of view, the observation will not contribute to the reduction of the state uncertainty. In other words, the entropy $H(x) = H(x|o)$ and the information gain will be zero. This way, by maximizing the mutual information gain, the cameras will effectively be in configurations in which more targets are present in the field of view.

The aforementioned problem boils down to how to compute $H(x|o)$ before making an actual observation. This entropy can be computed by taking into account that the tracking is being made by using an EKF. In the Kalman Filter (and so in the Extended Kalman Filter), the assumption is that the state distribution follows a Gaussian distribution with the mean x (the state estimate) and covariance P . Under this property, $H(x)$ and $H(x|o)$ are both differential entropies of Gaussian distributed variables, given by

$$H(x) = \frac{k}{2}(1 + \log(2\pi)) + \frac{1}{2}\log(|\Sigma|), \quad (2)$$

where k is the dimension of x and Σ is the covariance matrix of $p(x)$.

The first result consists in that the entropy depends only on the covariance matrix and thus the problem of computing $H(x)$ and $H(x|o)$ can be reduced to computing the covariance of the state estimate P_k and the state estimate after the observation P_{k+1} .

The equations of the Extended Kalman Filter show that the observation z_k is only incorporated in the innovation equation and later used in the state update equation. Therefore, the update covariance matrix $P_{k|k}$ can be computed prior to any observation. The same applies to the conditional entropy $H(x|o)$, as well.

In the most general case, the optimization can be achieved by an exhaustive search procedure over the parameter space of the camera. However, the structure of the sensor can be considered in order to find better optimization strategies.

III. SYSTEM OVERVIEW

The surveillance system consists of a set of active pan-tilt-zoom cameras which acquire images to be processed by a central controller. The controller is responsible for image processing tasks and for keeping track of each of the targets moving in the environment.

Each camera feeds the controller with new image frames and its corresponding state. The controller is responsible for processing the image and extracting target observations by fusing the available information and updating the targets state estimates (see Fig 1).

The result of the update is then fed to the decision controller, which is responsible for computing a new control policy for each camera.

The detection of new targets can be made both by carefully placed static cameras, or by the same active cameras used in the tracking process. The latter approach is the one followed in [11], however, this topic will not be addressed in this work.

IV. DETECTION AND TRACKING

The detection and tracking tasks are done using the CAMSHIFT (Continuously Adaptive Mean Shift) providing observations to an array of Extended Kalman Filters. It was originally developed for face tracking in vision based user interfaces, is

based on the mean-shift algorithm, which consists on a non-parametric technique for finding modes in probability distributions [4]. However, unlike the Mean-shift algorithm, the modified algorithm dynamically changes its color probability distributions.

In image processing tasks, the Mean-shift algorithm is used to find regions in images with similar color histograms. In tracking tasks, one has to store a target color model and that color model is used to track the same target in different frames.

The algorithm starts by building a color model for the object to track, which assigns a probability value to each pixel value. Then, for each image frame to be processed, the image is back projected with the pre-computed color model. The resulting image is a 2-dimensional probability distribution image encoding, at each pixel, the probability of a given pixel to belong to the original target. The final step is to assigned a fixed size window inside the window and compute the mean value of the probability distribution inside that window, which will correspond to the next window location. The process repeats until a convergence condition is met.

This results in a methodology robust to both lighting changes (by the use of a brightness invariant color model) and scale invariant, since the window size changes over time.

The observations are matched against existing (or new) EKFs and an update step is made for each observation available. The motion model is dependent on the target being tracked (the simpler approach is to use a random walk motion model) and the observation model the pinhole camera model.

V. SCHEDULING CAMERAS

Using the Information Theoretic approach proposed in [11], the camera scheduling problem is modeled as an information gathering problem. For each camera, a choice on the pan-tilt-zoom parameters is made based on the information gain in the state distribution of all targets given possible observations. The optimal pan-tilt-zoom configuration is the one which maximizes the information gain $I(x; o)$ between the state estimate and the observation. In other words, it is the one that leads to a greater increase in the certainty of the state estimate.

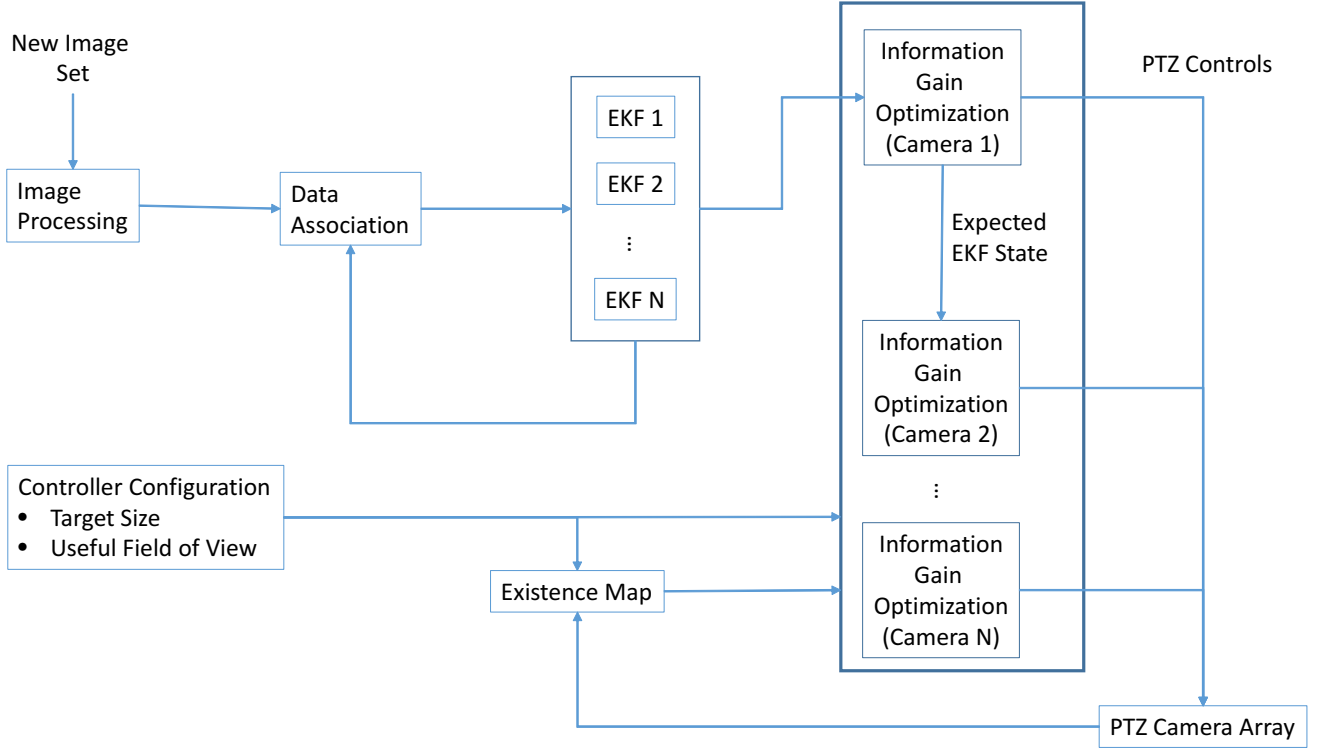


Fig. 1. System architecture

A. Single Camera Scenario

In a single camera scenario, a single camera is responsible for tracking and maintaining the trajectories of various targets which move freely across the environment.

The decision on which pan-tilt-zoom command to send to each camera is done by maximizing the mutual information gain $I(x_t; o_t)$ for all targets, with x_t being the target state estimate at time t from the EKFs and o_t the observation made.

Using the definitions of mutual information gain $I(x; o)$ and conditional entropy $H(x|o)$ [5], the cost function can be expanded

$$a^* = \arg \max_a I(x; o) = \arg \max_a H_a(x) - H_a(x|o) \quad (3)$$

$$= \arg \max_a H_a(x) + \int p(o) \int p(x|o) \log p(x|o) \partial x \partial o \quad (4)$$

$$= \arg \max_a H_a(x) + \int_v p(o) \partial o \int p(x|o) \log p(x|o) \partial x \\ + \int_{\bar{v}} p(o) \partial o \int p(x|o) \log p(x|o) \partial x, \quad (5)$$

where distribution $p(o)$ denotes the probability of making an observation of the target. Its domain of integration can be divided into two subdomains: v , which denotes all the camera configurations in which the target is visible and \bar{v} , which represents all the configurations in which the target is not visible.

Without further assumptions, this problem is hard to solve. However, recalling the EKF structure, predict and update steps, one obtains the following properties: (i) The probability distribution $p(x|o)$ corresponds to the state distribution after the update step. In other words, $p(x|o) \sim \mathcal{N}(x_{k|k}, P_{k|k})$. (ii) All state variables are gaussian distributed. This means all differential entropies can be computed in closed form according to equation (2). (iii) When there is no observation of the target, the update step is skipped. This means that the state is independent from the observation and $p(x|o) = p(x) \sim \mathcal{N}(x_{k|k-1}, P_{k|k-1})$. Applying these properties into equation (5), the objective function is simplified:

$$a^* = \arg \max_a H_a(x) + \omega(a)H(x^+) \\ + (1 - \omega(a))H(x^-), \quad (6)$$

where

$$\omega(a) = \int_v p(o) do \quad (7)$$

represents the observation function, which depends on the action a .

Note that all the state variables are Gaussian distributed and, by definition, the differential entropy for Gaussian distributed variables depends only on the variable covariance matrix. By observing the EKF equations, it can be seen that the observation o_k is only used in the innovation equation and in the state update equation. This enables the computation of P^+ without having an observation. Using the definition of differential entropy for a multivariate gaussian distribution (2) with the covariance update equation in the Extended Kalman filter, the cost functional can be simplified

$$a^* = \arg \min_a \omega(a) (\log |P^+| - \log |P^-|) \quad (8)$$

$$= \arg \min_a \omega(a) \log(I - K_k H_k). \quad (9)$$

The choice of the optimal configuration is made by optimizing the sum of the mutual information gains $I(x; o)$ for all targets. Despite the elegance of this cost function, in the general case its evaluation is intractable because it requires the exhaustive search over the configuration space of the camera. For each pan-tilt-zoom configuration of the camera, the observation model must be linearized anew and an EKF update must be performed to obtain the new Jacobian matrix H_k and K_k . In order to overcome that problem, in this work each target is modeled as a circle projected into the ground plane. The visible region v of each camera, in the camera coordinates, is also modeled as an ellipse around the center of the image. The term $w(a)$ in equation (9) is computed by projecting the target ellipse onto the image plane and computing its intersection with the visible region v of the camera. As it is shown in section VI, this will transform the cost function shape and make it suitable to use with numeric optimization procedures.

B. Multiple Camera Scenario

In a multiple camera scenario, multiple cameras can have observations of the same target. The incorporation of these observations is done using a sequential Kalman Filter. In this variation, a single

prediction step is made and each observation is used to make an update so that the update from the observation k uses the estimated position and covariance of the update from the observation $k-1$.

Each observation contributes to the covariance matrix with a factor of $(I - K_c H_c)$, where K_c is the Kalman gain from the observation of camera c and H_c is the Jacobian of the observation model for camera c .

The mutual information gain of a target for multiple observations

$$I(x; o_1, \dots, o_C) \propto \sum_{c \in C} \log |I - \omega(a) K_c H_c|, \quad (10)$$

is then obtained by combining equation (8) with the new covariance matrix update

$$P^+ = \left(\prod_{c \in C} (I - K_c H_c) \right) P^- . \quad (11)$$

C. Observation function

The observation function $\omega(a)$, equation (9), is a central component of the target tracking methodology as it effects on the convexity of the information gain.

The observation function can be just a flag indicating whether a point representing the target location, in world coordinates, is visible or not in the image:

$$\omega(a) = \begin{cases} 1, & \mathcal{P}(X_{gnd}; a) \in E_{img} \\ 0, & \text{otherwise} \end{cases} \quad (12)$$

where X_{gnd} is a point in the ground plane representing the target position, $\mathcal{P}(\cdot; a)$ is the projection operator on a set of points, which is defined by action a , i.e. the projection matrix $P_{3 \times 4}$ is modified by action a and E_{img} is the ellipse in the image plane concentric with the image rectangle. However, this observation function does not perform a smooth regularization of the cost function, making difficult to design iterative optimization algorithms.

By changing the model of the target from a single point in ground plane into a surface in the ground plane, one obtains a smoother optimization function. Modeling the shape of the target as a rectangle, taking into account the visible part normalized by the area of the imaged shape not truncated by the

field of view of the camera, the observation function takes the form

$$\omega(a) = \frac{A(\mathcal{P}(R_{gnd}; a) \cap E_{img})}{A(\mathcal{P}(R_{gnd}; a))} \quad (13)$$

where $A(\cdot)$ indicates area of a convex hull of points and R_{gnd} is a rectangle in the ground plane. Alternatively, modeling the target as an ellipse in the ground plane, the function becomes

$$\omega(a) = \frac{A(\mathcal{P}(E_{gnd}; a) \cap E_{img})}{A(\mathcal{P}(E_{gnd}; a))} \quad (14)$$

where E_{gnd} is an ellipsis in the ground plane.

VI. EXPERIMENTAL RESULTS

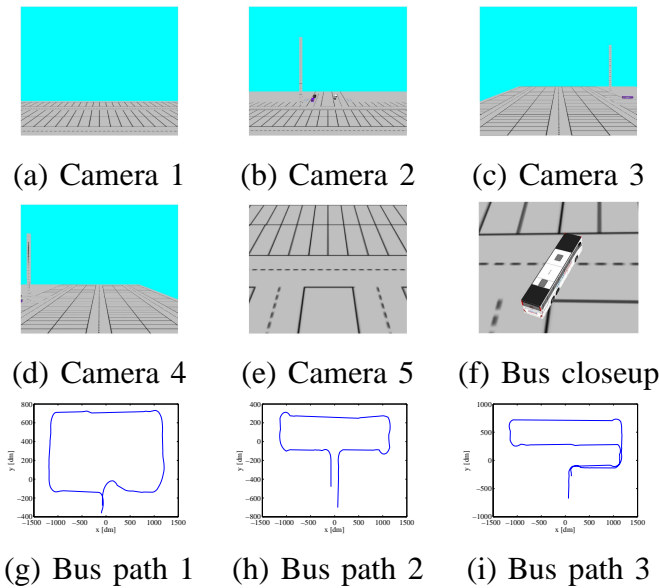


Fig. 3. Field of view of the cameras at rest orientation and typical bus trajectories.

The experiments described in this section are based in a virtual reality environment simulating a parking lot (see Fig. 3(a-f)). A number of buses, with the dimension of 2x10x2 m, cross the scene according to predefined trajectories shown in Fig. 3(g-i). These trajectories are generated using the car model $\dot{x} = \cos(\theta)\sin(\theta)V$, $\dot{y} = \sin(\theta)\cos(\theta)V$, $\dot{\theta} = \frac{\sin(\phi)}{L}V$ and $\dot{\phi} = \omega_s$, where V is the linear velocity of the bus and ϕ is the steering angle, both set using a joystick interface. Tracking is performed using five pan-tilt-zoom cameras located at fixed positions on the sides of the parking lot and in the entrance.

In order to evaluate the geometry of the cost function using different observation functions, the

cost function was evaluated from a set of values from the pan-tilt-zoom space for a camera. This was achieved by placing a target in a fixed location in front of the camera and making an observation. The camera model and its Jacobian were computed for the new pan-tilt-zoom configuration and, along with the new observation, used to update an existing EKF. The resulting Kalman gain and the linearized observation model were obtained and used in the cost function (9).

Figures 2(a) show the cost function evolution with both the zoom level and the pan-tilt values using the observation function (12). Figures 2(b) show the same evolution as in the previous case, but now using the observation function (14). The term $w(a)$ defined as in (9) changes the cost minimum from extreme values of zoom, pan and/or tilt to within-range, central, values.

The second set of figures was generated in a similar way from the former ones, however sampling the whole camera parameters space. Figure 4 shows slices of the cost function in the camera parameter space (pan-tilt-zoom) using different observation functions as described in Section V-C. Each row represents a fixed zoom configuration, with lower rows representing configurations with higher field of view (less zoom).

In column (a) of Fig. 4, is used $w(a)$ defined in (9). In column (b) is considered the same $w(a)$ however, in this experiment, there are two targets in the scene in different locations and with different state estimate covariances. In columns (c) and (d) the term $w(a)$ is the one defined by (13) and (14), respectively.

Figure 4(a) shows that the cost function is generally lower in configurations which make the target appear in the image plane with higher resolution, that is, when the camera is at its highest zoom or when the camera has the target of interest in the corner of the image. This result comes from the influence of the observation model (in particular, its Jacobian) used in the EKF in the cost function. While this is a good result in the sense resolution is maximized, uncontrolled zoom onto the target is not the desired behavior, since a minimum movement can make the target disappear from the camera field of view. Term $w(a)$, as defined in (13) or (14), acts a regularizing factor, managing the optimal zoom level to observe the target at constant zoom level and preventing observations just using the "corner

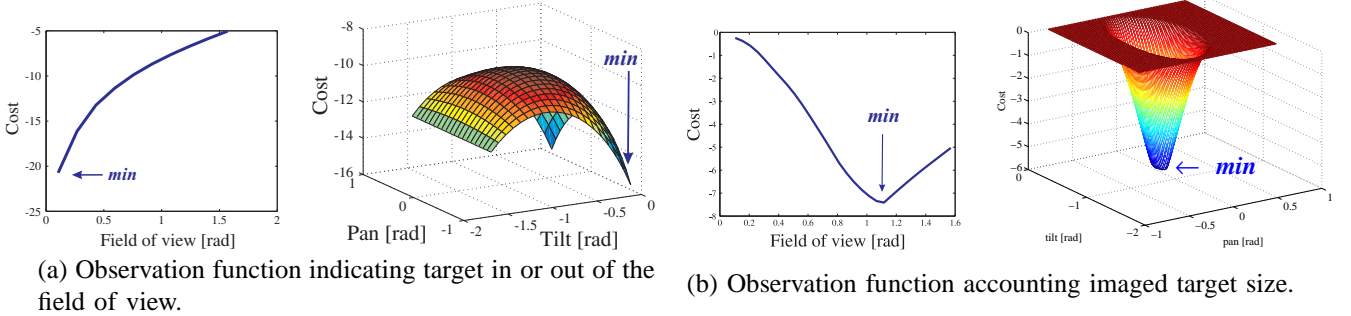


Fig. 2. Effect of the observation function into the cost function. In case (a) the observation function $\omega(a)$ is defined by (12) and therefore just indicates the target in or out the field of view. In case (b) $\omega(a)$ is defined by (14) and thus takes in account the imaged size of the target and the field of view. Two plots for each one of the cases, cost vs zoom (left) and cost vs pan and tilt (right).

of the eye”’.

In the last experiment four cameras are active and four targets (buses) enter the scene sequentially, separated by approximately 10sec, following different trajectories. The proposed scheduling methodology is tested with two alternative observation functions, (13) or (14). The pan, tilt and zoom parameters are estimated for all cameras in a round-robin manner (Fig. 1). Each search of the parameters is based in Nelder-Mead simplex direct search and is limited in the number of iterations. The optimization start points correspond to configurations in which the camera is pointing to a target or to the middle point of groups of targets, at the minimum zoom level possible. This is to enable the search of configurations in which multiple targets are inside the camera field of view.

Figures 6 (c) show a scenario in which the optimal decision is to observe multiple buses. As in the previous sample, the zoom level is also adjusted in order to keep both buses inside the camera field of view. Moreover, these results also show that the decision of which buses to cluster inside the same field of view is independent from the distance of the buses to the camera. Unlike other scheduling approaches, which take into account the distance of the targets to the camera in order to perform clusters of targets inside the field of view, this approach automatically takes into account the camera perspective effect. While on figure 6 (b) the buses are at the same distance to the camera are they are close together, in figure 6 (c) both buses are far from each other and from the camera.

Figure 7 shows a sample of the image frames acquired by the camera using both strategies As it is shown, modeling the targets as rectangles in




Fig. 6. PTZ imaging showing automatic clustering of multiple buses in the same field of view. Frame (a) shows automatic zoom selection when focusing a single target, while in (b) both buses are close together. In (c) both buses are positioned far away, clustering is due to the projective effect of the camera.

the ground plane (as described in equation 13) and using the proposed optimization procedure, the optimal decision is to observe the bus from the corners of the image. This shows that the observation function 13 is not effective enough to smooth the cost function. Although a formal proof could not be made, this behavior was not observed when modeling the targets as ellipses in the ground plane, as described in equation 14 (see figure 7 (b)).

Figure 5 show the evolution of the bus uncertainty over time. Results show that (14) allows for lower and smoother in time localization uncertainty in presence of distractors, as the other buses entering the fields of view of the cameras, due to its smoother




Fig. 4. Cost function slices at fixed field-of-view (zoom). Each row represents a different zoom level (field of view, top row $6[\text{deg}] \approx 0.1$ [rad], bottom row $75[\text{deg}] \approx 1.3$ [rad]). Colder colors represent lower cost functions. Configurations in which the target was not visible were assigned high cost. In cases (a,c,d) one camera observes one target. In case (b) one camera observes two targets. The observation function $\omega(a)$ is defined by (12) in cases (a) and (b), and is defined by (13) or (14) in cases (c) or (d), respectively.

nature allowing for more effective searches of the pan, tilt and zoom parameters.

First thing to notice in all the strategies is the intermittent behavior of the cameras. Sequence of

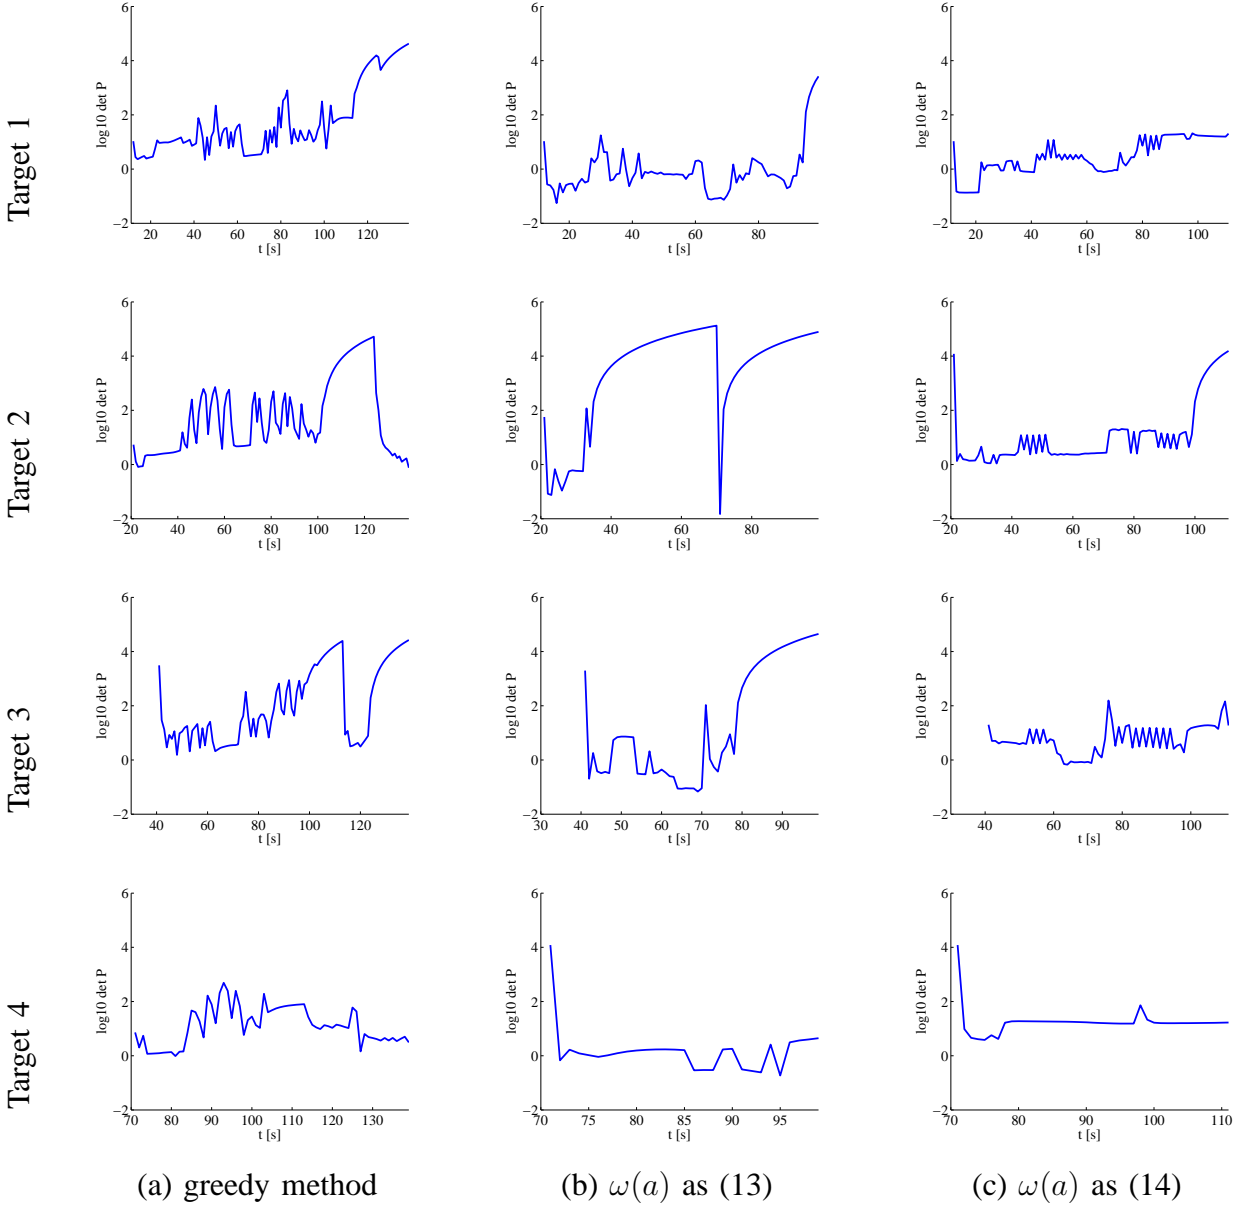


Fig. 5. Uncertainty of targets location using three different strategies for camera control. Each row corresponds to a different target. In (a) each camera sequentially chooses the target with largest entropy. In (b) and (c) the information gain is maximized with different observation functions. In the former, it is used an ellipse in the ground plane to model the targets (as described in equation 14), while in the later a rectangle is used (as described 13).




Fig. 7. Frame detail using two different observation functions at the same time instant. On the left (a), the optimization procedure gets trapped in a local minimum using the observation function 13. The same doesn't happen in (b), in which it is used the observation function 14

peaks mean that one or more cameras are continuously switching between targets. Despite undesired, this is to be expected since there is not a regularization term in the cost function in order to impose a finite speed limit for each camera.

The second point to address is the comparison between the strategies (b) and (c), using different observation functions. The first phenomenon to observe when using as observation function as in case (b) is not directly observable in figure 5. Figure 7 shows a sample of the image frames acquired by the

camera using both strategies. As it is shown, modeling the targets as rectangles in the ground plane (as described in equation 13) and using the proposed optimization procedure, the optimal decision is to observe the bus from the corners of the image. This shows that the observation function 13 is not effective enough to smooth the cost function. Although a formal proof could not be made, this behavior was not observed when modeling the targets as ellipses in the ground plane, as described in equation 14 (see figure 7 (b)).

VII. CONCLUSIONS AND FUTURE WORK

The results show the influence of the target modeling in the cost function proposed by [11]. By modeling the buses as ellipses projected onto the ground plane and the visible region of the camera as an ellipse in the image plane, the cost functions gain a defined structure, which simplifies the design of fast algorithms to search for the optimal policy.

When a single target is present inside the range of the active camera, the optimal command is to zoom on the target, according to the cost function. The regulator term controls the optimal zoom on the target, independently of the target location.

In a multi-target scenario, the overlapping cost functions can make the optimal command to keep both targets inside the camera field of view.

As future work, topics such as the application of the proposed algorithms in real systems need to be tackled, starting from the creation of a datasets likely to represent PTZ cameras.

REFERENCES

- [1] M Sanjeev Arulampalam, Simon Maskell, Neil Gordon, Tim Clapp, and Sanjeev Arulampalam. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *Signal Processing, IEEE Trans. on*, 50(2):174–188, 2002.
- [2] Horesh Ben Shitrit, Jerome Berclaz, Francois Fleuret, and Pascal Fua. Tracking multiple people under global appearance constraints. In *IEEE Int. Conf. on Computer Vision (ICCV)*, pages 137–144. IEEE, 2011.
- [3] Jerome Berclaz, Francois Fleuret, Engin Turetken, and Pascal Fua. Multiple object tracking using k-shortest paths optimization. *Pattern Analysis and Machine Intelligence, IEEE Trans. on*, 33(9):1806–1819, 2011.
- [4] Gary R Bradski. Computer vision face tracking for use in a perceptual user interface. 1998.
- [5] Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- [6] Joachim Denzler, Matthias Zobel, and Heinrich Niemann. Information theoretic focal length selection for real-time active 3d object tracking. In *IEEE Int. Conf. on Computer Vision (ICCV)*, pages 400–407. IEEE, 2003.
- [7] Francois Fleuret, Jerome Berclaz, Richard Lengagne, and Pascal Fua. Multicamera people tracking with a probabilistic occupancy map. *Pattern Analysis and Machine Intelligence, IEEE Trans. on*, 30(2):267–282, 2008.
- [8] Pedro Mendes Jorge, Arnaldo J Abrantes, and Jorge S Marques. Automatic Tracking of Multiple Pedestrians with Group Formation and Occlusions. In *VIIIP*, pages 613–618, 2001.
- [9] Keqin Liu and Qing Zhao. Indexability of restless bandit problems and optimality of whittle index for dynamic multichannel access. *IEEE Trans. on Information Theory*, 56(11):5547–5567, 2010.
- [10] Jerome Le Ny, Munther Dahleh, and Eric Feron. Multi-uav dynamic routing with partial observations using restless bandit allocation indices. In *American Control Conf.*, pages 4220–4225. IEEE, 2008.
- [11] Eric Sommerlade and Ian Reid. Probabilistic surveillance with multiple active cameras. In *IEEE Int. Conf. on Robotics and Automation (ICRA)*, pages 440–445. IEEE, 2010.
- [12] Eric Sommerlade, Ian Reid, et al. Cooperative surveillance of multiple targets using mutual information. In *Workshop on Multi-camera and Multi-modal Sensor Fusion Algorithms and Applications-M2SFA2 2008*, 2008.
- [13] Wiktor Starzyk and Faisal Z. Qureshi. Multi-tasking smart cameras for intelligent video surveillance systems. *8th IEEE Int. Conf. on Advanced Video and Signal Based Surveillance (AVSS)*, pages 154–159, August 2011.
- [14] Peter Whittle. Restless bandits: Activity allocation in a changing world. *Journal of applied probability*, pages 287–298, 1988.