Available at
www.ElsevierComputerScience.com
POWERED BY SCIENCE DIRECT®

Signal Processing 85 (2005) 607–624

SIGNAL PROCESSING

www.elsevier.com/locate/sigpro

Sensor management using an active sensing approach ☆

Chris Kreucher[a,b,*], Keith Kastella[b], Alfred O. Hero III[a]

[a]Department of Electrical Engineering and Computer Science, The University of Michigan, 1301 Beal Avenue, Ann Arbor, MI 48109-2122, USA

[b]General Dynamics Advanced Information Systems, P.O. Box 134008, Ann Arbor, MI 48113-4008, USA

Received 30 July 2004; received in revised form 11 November 2004

Abstract

An approach that is common in the machine learning literature, known as active sensing, is applied to provide a method for managing agile sensors in a dynamic environment. We adopt an active sensing approach to scheduling sensors for multiple target tracking applications that combines particle filtering, predictive density estimation, and relative entropy maximization. Specifically, the goal of the system is to learn the number and states of a group of moving targets occupying a surveillance region. At each time step, the system computes a sensing action to take, based on an entropy measure called the Rényi divergence. After the measurement is made, the system updates its probability density on the number and states of the targets. This procedure repeats at each time where a sensor is available for use. The algorithms developed here extend standard active sensing methodology to dynamically evolving objects and continuous state spaces of high dimension. It is shown using simulated measurements on real recorded target trajectories that this method of sensor management yields more than a ten fold gain in sensor efficiency when compared to periodic scanning.

Keywords: Sensor management; Machine learning; Active sensing; Multitarget tracking; Particle filtering; Joint multitarget probability density

1. Introduction

The problem of sensor management is to determine the best way to task a sensor or group of sensors when each sensor may have many modes and search patterns. Typically, the sensors are used to gain information about the kinematic state (e.g. position and velocity) and identification of a group of targets. Applications of sensor

☆This work was supported under the United States Air Force contract F33615-02-C-1199, AFRL contract SPO900-96-D-0080 and by ARO-DARPA MURI Grant DAAD19-02-1-0262. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force.

*Corresponding author. General Dynamics Advanced Information Systems, P.O. Box 134008, Ann Arbor, MI 48113-4008, USA. Tel.: +1 734 994 1200 X2717; fax: +1 734 994 5095.

E-mail addresses: ckreuche@umich.edu (C. Kreucher), keith.kastella@gd-ais.com (K. Kastella), hero@eecs.umich.edu (A.O. Hero III).

0165-1684/$ - see front matter © 2004 Elsevier B.V. All rights reserved.
doi:10.1016/j.sigpro.2004.11.004

management are often military in nature [36], but also include things such as wireless networking [28] and robot path planning [29]. There are many objectives that the sensor manager may be tuned to meet, e.g. minimization of track loss, maximization of probability of target detection, minimization of track error/covariance, and maximization of identification accuracy. Each of these different objectives taken alone may lead to a different sensor allocation strategy [36,38].

Many researchers have approached the sensor scheduling problem with a Markov decision process (MDP) strategy. However, a complete long-term (non-myopic) scheduling solution suffers from combinatorial explosion when solving practical problems of even moderate size. Researchers have thus worked at approximate solution techniques. For example, Krishnamurthy [26,27] uses a multi-arm bandit formulation involving hidden Markov models. In [27], an optimal algorithm is formulated to track multiple targets with an electronically scanned array that has a single steerable beam. Since the optimal approach has prohibitive computational complexity, several suboptimal approximate methods are given and some simple numerical examples involving a small number of targets moving among a small number of discrete states are presented. Even with the proposed suboptimal solutions, the problem is still very challenging numerically. In [26], the problem is reversed, and a single target is observed from a collection of sensors. Again, approximate methods are formulated due to the intractability of the globally optimal solution. Bertsekas and Castanon [1] formulate heuristics for the solution of a stochastic scheduling problem corresponding to sensor scheduling. They implement a rollout algorithm based on their heuristics to approximate the stochastic dynamic programming algorithm. Additionally, Castanon [5,6] formulates the problem of classifying a large number of stationary objects with a multi-mode sensor based on a combination of stochastic dynamic programming and optimization techniques. Malhotra [32] proposes using reinforcement learning as an approximate approach to dynamic programming. Very recently, Hernandez et al. [12] have used posterior Cramer-Rao bounds [41] to control the measurement sequence in a setting similar to that studied here.

Others have proposed using information measures a means of sensor management. In the context of Bayesian estimation, a good measure of the quality of a sensing action is the reduction in entropy of the posterior distribution that is expected to be induced by the measurement. Therefore, information theoretic methodologies strive to take the sensing action that maximizes the expected gain in information. The possible sensing actions are enumerated, the expected gain for each measurement is calculated, and the action that yields the maximal expected gain is chosen. Hintz et al. [15,16] focus on using the expected change in Shannon entropy when tracking a single target moving in one-dimension with Kalman Filters. A related approach uses discrimination gain based on a measure of relative entropy, the Kullback–Leibler (KL) divergence. Schmaedeke and Kastella [42] use the KL divergence to determine optimal sensor-to-target tasking. Kastella [21,23] uses KL divergence to manage a sensor between tracking and identification mode in the multitarget scenario. Mahler [30,31] uses the KL divergence as a metric for optimal multisensor multitarget sensor allocation. Zhao [47] compares several approaches, including simple heuristics, entropy, and relative entropy (KL).

Information-based adaptivity measures such as mutual information (related to the KL divergence) and entropy reduction are a common learning metric that have been used in the machine learning literature in techniques with the names "active object recognition" [8], "active computer vision" [44], and "active sensing" [10]. These techniques are iterative procedures wherein the system has the ability to change sensor parameters to make the learning task easier. The ultimate goal is to learn something about the environment, e.g. the class of an object, the orientation of a robot's tool, robot location.

A specific example of the role of information theoretic measures in machine learning is the repeated interrogation of an object to determine the object class. Denzler et al. [8] study a situation in which a camera has many adjustable parameters, including focal length, pan and tilt angles,