



AFRL-OSR-VA-TR-2015-0017

VISION-BASED AUTONOMOUS SENSOR-TASKING IN UNCERTAIN ADVERSARIAL ENVIRONMENTS

Amit Surana
UNITED TECHNOLOGIES RESEARCH CENTER EAST HARTFORD CT

01/02/2015
Final Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/ RTA
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.				
1. REPORT DATE (DD-MM-YYYY) 08-08-2015		2. REPORT TYPE Final Technical Report		3. DATES COVERED (From - To) DDMMYYYYDDMMYYYY
4. TITLE AND SUBTITLE Vision-Based Autonomous Sensor-Tasking in Uncertain Adversarial Environments		5a. CONTRACT NUMBER FA9550-11-C-0100		
		5b. GRANT NUMBER		
		5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Amit Surana (suranaa@utrc.utc.com) Allen Tannenbaum (arobertan@gmail.com) Computer Science Department, Stony Brook University Stony Brook, NY 11794-4400		5d. PROJECT NUMBER		
		5e. TASK NUMBER		
		5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) United Technologies Research Center, MS129-15, 411 Silver Lane, East Hartford, CT 06108		8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR / RTA 875 North Randolph Street, Suit 325 Room 3112 Arlington, Virginia 22203-1768		10. SPONSOR/MONITOR'S ACRONYM(S) AFOSR/RSA		
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION / AVAILABILITY STATEMENT A: Approved for public release; distribution is unlimited				
13. SUPPLEMENTARY NOTES Dynamical systems, control theory, computer vision, machine learning and statistical techniques				
14. ABSTRACT The key objectives of the proposed research was to build theoretical and computational foundations for developing efficient, robust and adaptive data-driven exploitation techniques and tools for automatic activity analysis in surveillance applications, and to incorporate dynamical systems, control theory, computer vision, machine learning and statistical techniques, through analytical and numerical methods, in the design of surveillance systems. The first goal is aimed at technology transition by creating and transitioning to Air Force generic tools that reduce analyst workload, enhance analysts situational awareness, and increase analyst efficiency and effectiveness in discovering and forecasting potential anomalous activities, and exploring hypotheses about those potential anomalous activities. Current autonomous sensor networks generate vast amounts of data while monitoring complex uncertain environments, provide limited actionable information, and are limited by the required number of human analysts. The aim of the second goal is to leverage dynamical systems and control theory to further optimize machine vision based surveillance systems such that they enable long term activity forecasting and early anomalous event detection; provide desirable trade-off between false alarms and missed detection; and exhibit robust performance under varying environmental conditions and scene contexts. Current machine vision systems have limited ability to exploit context in the				
15. SUBJECT TERMS				
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U		
				19b. TELEPHONE NUMBER (include area code)

Vision-Based Autonomous Sensor-Tasking in Uncertain Adversarial Environments

CONTRACT FA9550-11-C-0100

FINAL REPORT

Amit Surana (suranaa@utrc.utc.com)

United Technologies Research Center,
MS129-15, 411 Silver Lane, East Hartford, CT 06108,
and

Allen Tannenbaum (arobertan@gmail.com)
Computer Science Department, Stony Brook University
Stony Brook, NY 11794-4400

Objectives

The key objectives of the proposed research was to build theoretical and computational foundations for developing efficient, robust and adaptive data-driven exploitation techniques and tools for automatic activity analysis in surveillance applications, and to incorporate dynamical systems, control theory, computer vision, machine learning and statistical techniques, through analytical and numerical methods, in the design of surveillance systems.

The first goal is aimed at technology transition by creating and transitioning to Air Force generic tools that reduce analyst workload, enhance analysts situational awareness, and increase analyst efficiency and effectiveness in discovering and forecasting potential anomalous activities, and exploring hypotheses about those potential anomalous activities. Current autonomous sensor networks generate vast amounts of data while monitoring complex uncertain environments, provide limited actionable information, and are limited by the required number of human analysts. The aim of the second goal is to leverage dynamical systems and control theory to further optimize machine vision based surveillance systems such that they enable long term activity forecasting and early anomalous event detection; provide desirable tradeoff between false alarms and missed detection; and exhibit robust performance under varying environmental conditions and scene contexts. Current machine vision systems have limited ability to exploit context in the activity analysis, forecast activities, and analyze complex scenes with multiple interacting entities. Specific applications include autonomous aerial surveillance systems that cover broad areas of military operations, camera security systems that cover large crowded areas in urban environments, and large-scale wireless sensor networks that must minimize power consumption while providing actionable system state information.

Summary of Accomplishments

We summarize below the main accomplishments of this research program. As pointed above, cross fertilization of concepts from dynamical systems, control theory, computer vision, ma-

chine learning and statistics provided the necessary theoretical and computational foundations for developing novel techniques for human activity modeling, learning, classification and forecasting. These techniques:

- can be used for activity modeling and analysis at multiple spatio-temporal scales i.e. microscopic, mesoscopic and macroscopic;
- incorporate temporal information in a form which not only enables early activity classification but also activity forecasting, and thus could increase analyst efficiency and effectiveness in exploring multiple hypothesis and early discovery of potential anomalous activities;
- more effectively exploit context in activity analysis which could potentially lead to lower rates of false alarm and missed detection, and thereby reduce analyst workload;
- could enhance analysts situational awareness about activities in the crowded and cluttered monitored space, and help analysts develop and maintain a comprehensive picture of the operational environment.

Our work can be broadly categorized into two main themes: single agent activity analysis, and crowd activity analysis. For single agent activity analysis we have taken a microscopic viewpoint relying on tracking individual agents in videos. For analyzing crowd behavior we considered a mesoscopic/macroscopic viewpoint which utilizes coarse level features extracted from videos, and does not rely on tracking individuals.

In the first theme of our work, for modeling and analysis of long-term goal-oriented single agent activities, we developed a hierarchy of increasingly complex statistical models. The details are presented in Section 1.1. In order to capture global motion patterns and detect anomalous behavior, we initially developed a Markov modeling approach, and used it in conjunction with geometric active contour based multi-target tracking and statistical change detection methods. While this approach succeeded in detecting anomalies in a complex video with multiple individuals, the Markov models were found not to be effective in forecasting behavior. To alleviate this limitation, we proposed a Markov Decision Processes (MDP) framework for goal-oriented activity modeling and analysis. MDP provides a more natural framework to capture rational human activities which can be thought of as being driven by immediate rewards, expected future rewards and goals. For learning MDP models from trajectory data, one could use standard techniques from Inverse Reinforcement Learning (IRL). However, applying the standard MDP/IRL framework in computer vision applications require several additional considerations such as noisy and unlabeled trajectory data, and non-stationary rewards which drive agent behaviors.

To address these challenges, we developed extensions of MDP/IRL framework which enables unsupervised learning from noisy trajectory data, and analysis of multi-scale switching behaviors. For this we introduced two new classes of MDP models: hidden variable MDPs (hMDPs) and switched MDPs (sMDPs), and developed advanced Markov Chain Monte Carlo (MCMC) learning techniques based on Bayesian Nonparametrics (BNP). Rather than comparing models that vary in complexity (and choose the best one) like in classical approaches, the BNP approach is to fit a single model that can adapt its complexity to grow as

more data is observed. This is essential in complex settings, where the space of models to be searched is difficult to efficiently enumerate and explore. We also developed online Bayesian techniques for behavior classification and forecasting with the different MDP model representations discussed above. We applied our MDP framework in a simulated urban environment, and demonstrated several desirable features, including: long term behavior prediction and early behavior classification; robustness to noise and better generalizability with limited training data; and ability to encapsulate behaviors in terms of scene features (and not scene locations), and thus providing basis for transfer learning.

In the second theme of our work, we developed new techniques for tractable modeling and analysis of crowd behaviors. The details are presented in Section 1.2. For crowded scenarios, microscopic viewpoint faces considerable difficulty in moderate to high density crowds as tracking performance can significantly deteriorate. Therefore, we resorted to mesoscopic/macroscopic modeling which tend to be more reliable. In our studies mesoscopic representation considers crowd as a collection of dynamically interacting and evolving groups of individuals, while macroscopic representation treats crowd as one global entity.

In this regard, we developed a variational framework for group detection and tracking in crowds. This framework is based on dynamic active contour driven by optical flow, and fuses temporal and intensity distribution information explicitly into a single framework. Dynamic active contours are used to spatiotemporally segment crowd and detect and track groups. A level set active contour formulation is used to account for global topological changes in the contour shape, i.e. splitting and merging. Optical flow is used to drive the dynamic contours. Furthermore, geometric observer theory can be used in conjunction with this framework for error correction leading to more robust detection and tracking. Our numerical experiments showed high group detection rate despite splitting, merging and collisions in complicated real world videos.

We also developed an approach for crowd anomaly detection based on system identification techniques. In this approach the video is represented at a macroscopic scale using a linear dynamic texture model. We utilized a subspace system identification method based on Hankel matrix to extract relevant dynamics of noisy low level features extracted from the video. The spectral properties of the Hankel matrix encode useful information about the underlying dynamics, and changes in those properties can be used to detect anomalous behaviors. In particular, we demonstrated that by monitoring rank of the Hankel matrix, we could robustly detect onset of panic in crowd videos. Furthermore, application of this approach to very dense crowd video scene revealed existence of very low order dynamics. Using this insight, we developed another macroscopic approach for dense crowd behavior analysis which treats crowds as a fluid flow driven by optical flow in the images. Given this analogy with fluid flow, we employed geometric, statistical and spectral concepts from nonlinear dynamical systems techniques to detect coherent motion patterns in such flow fields and use them for crowd motion segmentation and change detection in crowd behavior. In particular we investigated Finite Time Lyapunov Exponents, Perron Frobenius Operator and Koopman Operator based analysis, and used them for crowd segmentation and characterization of internal dynamics within the segments in several real world videos.

Furthermore, we extended our system identification/dynamical system framework discussed above for modeling of more general scenes. In particular we developed a novel nonlin-

ear approach for modeling of complex dynamic texture videos based on Koopman operator theoretic method. Koopman operator is linear but infinite dimensional operator, and captures full nonlinear behavior. We exploited this aspect to construct a linear stochastic system in Koopman mode space, and used it as a generative model for nonlinear dynamic textures. Through various complex texture videos, we showed superior modeling performance of our approach over other methods proposed in the literature. This Koopman based data driven model reduction technique is currently being transitioned at UTRC in context of other applications including rotorcraft prognostics and health management, and big data streaming analytics.

Outline of Report

In Chapter 1 we describe in more detail the key ideas of our technical approach and some numerical results. Full details can be found in the associated publications. In Chapter 2 we briefly outline how the techniques developed in this program are being transitioned at UTRC. Chapter 3 lists the UTRC personnel supported under this program, and finally Chapter 4 lists the publication which resulted from this contract.

Contents

1	Summary of Research Results	6
1.1	Single Agent Activity Analysis	6
1.1.1	Markov Models for Statistical Behavior Analysis	6
1.1.2	Markov Decision Processes for Goal Oriented Behavior Analysis . . .	9
1.1.3	Hidden Variable MDPs for IRL with Noisy Data	11
1.1.4	Switched MDPs for Multiscale Behavior Analysis	13
1.2	Crowd Activity Analysis	16
1.2.1	Variational Framework for Detecting and Tracking Groups in Crowds	16
1.2.2	Hankel Operator based Anomaly Detection	17
1.2.3	Nonlinear Dynamical System Analysis for Dense Crowd Segmentation	19
1.2.4	Koopman Operator based Nonlinear Dynamic Textures	21
2	Transitions at UTRC	23
2.1	Data Driven Nonlinear Model Reduction for PHM/Big Data Applications . .	23
3	Personnel Supported	24
4	Publications	25

Chapter 1

Summary of Research Results

Our research is concerned developing new techniques for behavior/activity modeling, learning, classification and forecasting in complex surveillance videos. In our work we are using motion patterns as an approach for activity modeling and analysis at multiple spatio-temporal scales, i.e. microscopic, mesoscopic and macroscopic.

Our work can be broadly categorized into two main themes: single agent activity analysis, and crowd activity analysis. For single agent activity analysis we have taken a microscopic viewpoint relying on tracking individual agents in videos. For analyzing crowd behavior we considered a mesoscopic/macroscopic viewpoint which utilizes coarse level features extracted from videos, and does not rely on tracking individuals. In order to evaluate and validate our techniques, we used three main sources of surveillance datasets: public domain video data, UTRC’s desktop agent based modeling and simulation environment, and UTRC’s Multi Camera Cafeteria testbed. The selected datasets are representative of complex surveillance scenarios in challenging environments, and are rich with multi-scale single agent and multi-agent activities, including threatening behaviors and correlated crowd behaviors.

1.1 Single Agent Activity Analysis

In this section we describe several classes of increasingly complex statistical models for activity analysis based on individual tracks. We first discuss Markov models which can be employed to capture global motion patterns and detect anomalous behavior. We next describe a Markov Decision Processes (MDP)/Inverse Reinforcement Learning (IRL) framework which provides a powerful paradigm for representing and analyzing long-term goal-oriented behavior. Furthermore, by utilizing the flexibility of a Bayesian Nonparametric framework in Bayesian IRL, we discuss extensions which enable unsupervised IRL with noisy trajectory data and analysis of multi-scale switching behaviors.

1.1.1 Markov Models for Statistical Behavior Analysis

Our Markov model based framework for statistical trajectory modeling and anomaly detection involves three main steps [c1]: 1) Obtaining individual object tracks in the scene; 2)

Choosing a set of coarse variables as a function of object tracks and building a Markov model which captures global motion patterns; and 3) Using coarse Markov models from step 2) in a change detection framework for detecting anomalous behavior.


We used a geometric particle filtering approach for multi-target tracking and obtaining object tracks. In this approach the particle filtering based tracking/recognition is augmented with knowledge of object shape to guide segmentation in uncertain regions. More specifically, the shapes and the deformations of the objects are modeled using geometric active contours, and a continuous state hidden Markov model (HMM) is defined whose state comprises the continuous contour and its velocity (which consists of its local and global deformations), while the image at a given time forms the observation. To speed up the particle filtering we use an approximate importance sampling density which requires only sampling the space of affine deformations while approximating the local deformation by the mode of its posterior [33, 27, 28]. In our implementation, we used an active contour model (with level set representation) whose evolution is based on the Bhattacharyya distance [19]. This model can be viewed as a generalization of those segmentation methods in which the active contours maximize the difference between a finite number of empirical moments of the distributions “inside” and “outside” the evolving contour. The model is very versatile and flexible since it allows one to easily accommodate a number of diverse image features. Furthermore, it can incorporate both local and global information, and extends naturally to multiple contour evolution. Incorporating prior shape knowledge in the curve evolution step is many times necessary when dealing with occlusions. We have however dealt with this issue by incorporating the shape information in the weighting step of particle filtering instead of the curve evolution step (see [28] for details).

Application of change detection directly to object tracks treated as output of the HMM is a computationally intensive process, especially when the change parameter is unknown [34, 14]. Moreover, the continuous HMM only captures local object motion (useful for tracking), and is thus not very useful for capturing global motion pattern which is needed to detect anomalies. We therefore propose to use coarse statistical models instead, which are derived by coarsening the space of object tracks based on feature variables which capture relevant aspects of the object’s global motion. Motivated by the Mori-Zwanzig-Shannon projective approach to modeling complex phenomena [17], we have used an empirical Markov model for coarse representation of object dynamics. Furthermore, the use of such coarse models facilitates the application of change detection methods. We propose to use a CUSUM-like change detection test statistic based on the Donsker-Varadhan rate function [17]. This statistic can be efficiently computed using a prior Markov model (learnt from historical data representing nominal behavior) and a real time calibrated Markov model obtained from the video data.

We have demonstrated our trajectory modeling approach described above in a challenging video with multiple pedestrians moving in a cluttered and occluded indoor UTRC cafeteria environment. Figure 1.1a shows the image plane trajectories of four individuals obtained by using the geometric particle filter. Despite clutter and occlusions, our geometric filter is able to maintain all the tracks. Note that these tracks have different spatiotemporal behavior, i.e., differ in how much time each individual spends near tables and overall motion pattern relative to the tables. In order to distinguish these behaviors, we use a coarse model derived



(a) Tracks



(b) Anomaly Scores

Figure 1.1: a) shows the tracks of four people, while b) shows the anomaly score. Despite clutter and occlusions, our geometric filter is able to maintain all the tracks. With our anomaly detection approach, the magenta track is easily picked as anomalous due to extended waiting time near the tables, while the red track begins to appear anomalous when it circles around the tables.

from the tracks in the physical space.

For learning a coarse nominal model we have developed an agent based model (ABM) simulator for the cafeteria. The cafeteria environment is represented in the form of a discretized elevation map, indicating the computational cells occupied by tables. We assume each agent is goal-oriented and attempts to follow the shortest path from an entrance to its destination cell near a table, spends some time there, and then follows a shortest path to one of the exits. Based on these simulated trajectories a single empirical Markov model of nominal behavior is constructed using two coarse variables: cell location and dwell time. Dwell time measures the period of activity (zero if there is motion) and inactivity (one for every consecutive time-step of no-motion) and allows us to capture memory in the dynamic process in an efficient manner [30]. In order to derive the coarse model from video trajectories, we project the tracks in the image plane onto the physical space using camera calibration parameters. Consequently, each track can be mapped into a sequence of cells defined in the ABM and the coarse variables are determined.

Figure 1.1b shows the anomaly score as a function of time, computed using the change detection statistic discussed earlier, for each tracked individual. The black line shows the chosen threshold. There is one-to-one correspondence in the colors used in Figures 1.1a and 1.1b. Clearly, the magenta track is easily picked as anomalous due to extended waiting time near the tables, while the red track begins to appear anomalous when it circles around the tables.

1.1.2 Markov Decision Processes for Goal Oriented Behavior Analysis

In this section we discuss the application of Markov Decision Processes (MDPs) based goal oriented behavior learning, classification, and prediction. MDPs have been a popular approach for modeling sequential decision making [25]. They offer several advantages for human behavioral modeling in surveillance scenarios [13]. Firstly, rational human activities can be thought of as being driven by immediate rewards, expected future rewards, and goals, which can be naturally captured as an MDP. In the MDP setting the agent’s motion trajectory, i.e., state-action pairs, is generated by executing an optimal policy based on agent preferences or rewards. Secondly, an MDP representation encapsulates behaviors in terms of physical scene features and not physical location, and so has the ability to generalize to novel scenes, enabling transfer learning.




Figure 1.2: Learning Markov Decision Processes using Inverse Reinforcement Learning. a) shows the reward basis function. b) right plot shows the trajectories sampled from the learned MDP which appear very similar to the training dataset shown in left plot.

In computer vision applications, one only has access to trajectory data extracted from videos which indirectly represents how agents behave in the given environment. To apply the MDP framework to represent this behavior, one needs to learn an agent’s rewards/preferences which drive the agent behavior based on the observed trajectories. Given the environment model, the problem of finding the reward function that explains the observed agent’s behavior is termed as Inverse Reinforcement Learning (IRL). The IRL problem has been addressed in the literature using two main formalisms: *reward learning* (i.e., determining reward parameters) and *apprenticeship learning* (i.e. direct policy learning). The IRL problem is inherently ill-posed, since there are infinitely many reward functions that may yield the policy as optimal [23, 5]. To address this non-uniqueness, different approaches have been proposed in literature to encode preferences over the reward or policy function spaces. These approaches can be broadly categorised into *Optimization based IRL* and *Bayesian IRL*. Optimization based IRL approaches encode preferences over the reward or policy function spaces by using appropriate objective functions and/or constraints [23, 29, 2, 32, 35, 21, 22]. Bayesian approaches formulate the reward preferences in the form of prior distribution and behavior




Figure 1.3: Prediction of agent future paths based on past observed behavior. The darker yellow shades imply higher chance of visiting those areas in future. Initially there is ambiguity in the final goal of the agent, but as more data is available the MDP approach is able to correctly predict in advance the final goal and the most likely future path towards it.

compatibility as a likelihood function [26, 5].

To illustrate the MDP representation for goal oriented behavior, we consider a simulated surveillance problem in an urban environment. In this scenario the goal-oriented agents move around to reach their desired destinations near the buildings, denoted by A, B and C in the Figure 1.2a. The urban environment is represented in form of the discretized elevation map indicating the computational cells occupied by the buildings. The agent's state is represented by the cell it occupies and there are 4 available actions: move *north*, *south*, *east* and *west*. The reward function for each agent is parameterized in terms of four basis functions: $\phi_i, i = 1, \dots, 4$, with ϕ_1 and ϕ_2 shown in the Figure 1.2a. The function ϕ_1 penalizes areas occupied with buildings which need to be avoided, while ϕ_2, ϕ_3, ϕ_4 have high values near the goal destinations A, B and C, respectively. Accordingly, we assume there are 3 types of agent behaviors: *BehA*, *BehB* and *BehC*. Here, *BehA* denotes the behavior where agent has preference for destination A, with similar interpretation of *BehB* and *BehC*.

Based on these reward preferences we compute optimal policies and generate near optimal trajectories for each type of behavior. The left sub-figure in Figure 1.2b shows the labeled training data for *BehB*. In order to learn an agent's reward preferences/policy from this training dataset, we used a linear programming approach for IRL which is based on a dual approach for solving for MDP policy [32]. The right sub-figure in Figure 1.2b shows the trajectories sampled from the learned MDP which appear very similar to the learning dataset. Note that to get similar performance with a Markov model, much more learning data would be required.

Once the MDP model has been learned it can be used for behavior prediction and classification for which we have developed an online Bayesian approach. Figure 1.3 shows the prediction of expected future possible paths (denoted by yellow) of an agent based on its observed past behavior (denoted by a black track). Initially there is ambiguity in the final goal of the agent, but as more data is available the MDP approach is able to predict well in advance the final goal and the most likely future path towards it. Markov models degenerate to a random walk when used for long term forecasting, and fail to clearly delineate an agent's future path [13]. Thus, by incorporating goal-oriented behavior using an MDP




Figure 1.4: Behavior classification in a convergence scenario with 6 agents with different goals. The agents shown by red/green track are correctly classified as moving towards B (and others as not having that goal) in advance of they actually reaching that location.

representation, one can achieve a greater predictive capability.

Application of the MDP framework for behavior classification is illustrated in Figure 1.4 where agents are moving towards their different goal locations. Here the goal B is deemed important, and the objective is to classify which agents are heading towards that goal. The right sub-figure in the Figure 1.4 shows the likelihood of different agents converging towards goal B. Clearly the agents shown by red/green track are correctly classified as moving towards B (and others as not having that goal) in advance of they actually reaching that location.

1.1.3 Hidden Variable MDPs for IRL with Noisy Data

Application of the MDP framework discussed above for trajectory-based activity analysis in computer vision applications requires additional considerations. Firstly, the trajectories which are output from a tracking algorithm are typically noisy. As a result, the true state (e.g. position) of agent is not directly observable. Secondly, in videos there are typically many agents, and the number of different behaviors and the behavior labels for each agent trajectory may not be known a priori, and should also be learned in addition to the rewards. We developed a new Bayesian IRL framework for unsupervised learning from noisy trajectory data [c2].

To deal with noisy data, we use a hidden variable MDP (hMDP) representation. In hMDP, observation uncertainty is modeled via a hidden state variable as in a Partially Observable Markov Decision Process (POMDP). However, hMDP is different than POMDP in the sense that the agent is not uncertain about its own state (and does not have to account for that uncertainty in making decisions), it's only the observer who has noisy observations of agent's state. For unsupervised learning with an hMDP representation we used a Bayesian IRL (BIRL) framework. We first developed hMDP BIRL (hBIRL) techniques, assuming noisy trajectory labels are given. For this we exploit that, for a fixed policy, hMDP reduces

to a Hidden Markov Model (HMM). Hence, Markov Chain Monte Carlo (MCMC) methods developed for parameter learning in HMMs, can be employed. In particular, we developed two approaches for hBIRL: one is based on likelihood recursion, which marginalizes over the hidden state sequence in the underlying HMM, and the other uses forward-backward Gibbs sampling. The latter approach is preferred as it leads to a faster mixing Markov chain.

We next extended the hBIRL framework to a nonparametric setting for which we employed a Dirichlet Processes (DP) mixture model as a prior over the behavior clusters, and use a MCMC sampling procedure. During this sampling, the clusters, reward parameters per cluster, and the underlying state sequence per trajectory are sampled sequentially utilizing a Chinese Restaurant Process representation of the DP mixture model [20]. This BNP approach automatically partitions the trajectories without the need to specify a priori the number of distinct behaviors present in the dataset.




Figure 1.5: a) Noisy training trajectories. Subplots b) and c) show posterior samples of reward parameter (red triangles are true reward parameters) from a MCMC run for two different noise levels. One can see 3 distinct clusters in reward parameter space and as expected, for the low noise case these clusters are more prominent, while for high noise case they become fuzzy. Subplot d) shows behavior classification for 6 agents for the high noise case: solid curves are the true trajectories, while dashed curves are the observed noisy trajectories. Agents 1 and 2 follow *BehB*, agents 4 and 5 exhibit *BehA*, and agents 3 and 5 move according to *BehC*. Clearly, agents 1 and 2 shown by red/green tracks are classified correctly as moving towards B, considerably in advance.

We demonstrated our nonparametric-hBIRL (NP-hBIRL) for unsupervised behavior learn-

ing in the simulated urban surveillance scenario discussed in the previous section. Figure 1.5a shows a subset of the unlabelled noisy training data. Figures 1.5b-c show posterior samples from a MCMC run for the two noise levels. One can see 3 distinct clusters in reward parameter space corresponding to *BehA*, *BehB* and *BehC*. As expected, for the low noise case these clusters are more prominent, while for high noise case they become fuzzy.

We also developed a recursive online Bayesian approach for behavior classification and prediction with the hMDP models. Figure 1.5d shows application to behavior classification problem. Given the noisy trajectories of the different agents which are updated over time, the goal is to classify which agents are most likely moving towards the critical destination B. The right sub-figure in 1.5d shows the likelihood of different agents heading towards B. Clearly, agents 1 and 2 shown by red/green tracks are classified correctly as moving towards B, considerably in advance of when they actually reach that destination, despite the high noise in the observed trajectories.

1.1.4 Switched MDPs for Multiscale Behavior Analysis

IRL approaches proposed in the literature (as discussed above) typically assume that the agent’s behavior can be described by a single underlying reward function. However, human decision making routinely involves choosing and/or switching among temporally extended courses of action over a broad range of time scales. We have developed a switched MDP (sMDP) modeling framework to capture temporally extended courses of action, and develop a BNP approach to learn such models from the behavior data [c3].

sMDP consists of a finite number of modes each modeled by a MDP representing a simpler behavior, and a Markov switching process which selects a sequence of modes over time. The proposed sMDP framework is along the lines of different generalizations of MDP models which have been put forth for representing complex multi-scale temporal human behavior [31]. Another motivation of using sMDP models comes from the success of using switched linear dynamical systems to explain complex nonlinear behaviors in a variety of real-world applications [7]. In sMDP the number of Markov modes is typically unknown a priori and should also be learned in addition to reward preferences in each MDP mode. We take a BNP approach for defining a prior on the sMDP model parameter space. Specifically, we use a Sticky Hierarchical Dirichlet Process (sticky HDP) introduced in [7], as the prior. The sticky HDP model better captures the temporal mode persistence representing a temporally extended course of action, and thus provides more control over the number of hidden modes that are inferred. We have developed an efficient inference algorithm based on MCMC sampling to obtain posterior samples of the sMDP model parameters given the data. This BNP IRL approach makes fewer assumptions about the underlying dynamics than are required by parametric ones, allowing the data to drive the complexity of the inferred model [9].

We compare the performance of our BNP IRL approach with BIRL (see [26]) in learning temporally switching behavior from the training data, a subset of which is shown in 1.6a. Figures 1.6c-d show the posterior samples for reward parameters for one of the MCMC trials (with 5000 iterations) using the two methods, respectively. BNP IRL identifies 3 clusters corresponding to *BehA*, *BehB* and *BehC* which compose *BehABC*. *BehABC* here denotes a behavior where agent switches between *BehA*, *BehB* and *BehC*, with very high probability




Figure 1.6: a) Sample trajectories from the training dataset with switching behavior in an urban like environment. The color of the trajectory segments correspond to the destination of the same color, the symbol $\triangle$ denotes the starting point of trajectories, and $\square$ denotes the end points. b) Posterior samples from one of the BNP IRL MCMC trials. BNP IRL identifies 3 clusters corresponding to *BehA*, *BehB* and *BehC* which compose *BehABC*. c) Posterior samples from one of the BIRL MCMC trials. In this case, MCMC samples are randomly distributed in reward parameter space as no combination of parameters explains the data well, and the method gets stuck at a random point in parameter space where most moves are equally poor. d) Hamming distance averaged over multiple MCMC trials (left plot), and comparison of averaged log-likelihood for the two methods (right plot). The log-likelihood function has significantly lower values and does not improve much over iterations for BIRL when compared to BNP IRL, illustrating that the sMDP learned using BNP IRL can explain the data more effectively than a single reward MDP learned using BIRL.

of following *BehB* once that mode has been chosen. Fig. 1.6d shows the convergence of Hamming distance (which represents the error between true modes and inferred modes [7]) to low values. On the other hand, using standard BIRL we found that: 1) MCMC samples are randomly distributed in reward parameter space (see fig. 1.6c) as no combination of parameters explains the data well, and 2) the method gets stuck at a random point in parameter space where most moves are equally poor. This is clearly illustrated in figure




Figure 1.7: Behavior prediction: top row and bottom row figures use models learned using BNP IRL and BIRL, respectively. As more trajectory is observed, our BNP IRL method is able to correctly classify behavior to be of type *BehABC* in advance, and thus accurately predict the agent’s future behavior. For BIRL case the behavior posterior switches between different behaviors and fails to correctly predict the future path.

1.6d where we compare the log-likelihood function (for training data) averaged over multiple sets of MCMC trials. It can be seen that the log-likelihood function has significantly lower values and does not improve much over iterations for BIRL when compared to BNP IRL. This shows that the sMDP learned using BNP IRL can explain the data more effectively than a single reward MDP learned using BIRL [5].

We have also developed a Bayesian approach for classification and prediction of agent behavior represented by an sMDP model. This approach is online, in which the posterior on the behavior class, and prediction of average future behavior is updated based on agent’s behavior observed so far. We demonstrated this in a simulated surveillance scenario, where we show how an sMDP model representing temporally rich behavior can be used for more effective behavior classification and prediction, compared to the standard MDP model learned using classical BIRL approach. This comparison is shown in Figure 1.7 where we consider the problem of predicting behavior of an agent who is following a deceptive behavior *BehABC* (denoted by a black track in fig. 1.7). Figures 1.7a-c show the expected future occupancy map (denoted by yellow) based on models learned using the BNP IRL approach, and the figure 1.7d shows the behavior posterior as a function of time. As more trajectory is observed, our method is able to correctly classify behavior to be of type *BehABC* in advance, and thus accurately predict the agent’s future behavior. Figures 1.7 e-g show similar results, but with models learned using BIRL. In this case, the behavior posterior switches between different behaviors as shown in the figure 1.7h, and so does the future predicted path. The single reward MDP is not effective in predicting temporally switching behavior.

1.2 Crowd Activity Analysis

In the second theme of our work, we developed new techniques for tractable modeling and analysis of crowd behaviors. The conventional bottom up approach treats crowd as a collection of individuals, and thus relies on individual detection and tracking to analyze behaviors. This approach faces considerable difficulty in moderate to high density crowds as tracking performance can significantly deteriorate. For such scenarios mesoscopic representation which considers crowd as a collection of dynamically interacting and evolving groups, or macroscopic representation which treats crowd as a one global entity tend to be more reliable. In this regard, we developed several techniques for mesoscopic and macroscopic crowd analysis. These techniques are based on concepts from system identification, nonlinear dynamical systems, and variational formulation utilizing dynamic active contours.

1.2.1 Variational Framework for Detecting and Tracking Groups in Crowds

In this section we describe a mesoscopic approach for detecting and tracking groups in crowds. There is a sociological hypothesis that the majority of people in the crowd cluster in small groups. Finding small groups traveling together is thus a fundamental problem in understanding crowds, and improving situation awareness and emergency response during public disturbances.

For detecting and tracking groups in crowds, we have developed a variational framework based on dynamic active contours in conjunction with optimal mass transport based optical flow [c4]. This allows us to fuse temporal and intensity distribution information explicitly into a single framework. The key idea is to use optical flow as the macroscopic model of crowd motion, and use that to drive dynamic active contours which spatiotemporally segment crowd into groups of individuals. The main advantages of our approach are as follows. Firstly, the use of level set formulation with dynamic active contour enables frame to frame tracking of groups in crowds accounting for global topological changes including merging and splitting of contours. Secondly, the optical flow we use to drive the dynamic contours is based on optimal mass transport. The dynamic textures typical of crowd videos possess intrinsic dynamics and so cannot be reliably captured by the standard optical flow methods as used in several previous studies. Thirdly, our variational framework can be readily extended to incorporate richer crowd motion models and employ geometric observer theory [24] for more robust group detection and tracking. Using the key ingredients of the above general framework, we have also developed a simplified variational approach for group motion detection.

Our numerical experiments show high detection rate of macro crowd behaviors such as splitting, merging, and collisions in complicated real world videos taken from event recognition videos from the 2009 PETS benchmark dataset. Figure 1.8 shows some frames with identified groups and their dominant motion for the splitting scenario. In Figure 1.9 we show the results of the merging scenario, while Figure 1.10 shows the identified groups and their dominant motion for the colliding and merging scenario. We are currently developing appropriate group motion models to incorporate them in the above variational framework [p2].

This is expected to not only improve the tracking performance, but also enable prediction of events such as splitting, merging, panic etc.



Figure 1.8: Detecting and tracking group splitting behavior. The level set formulation used here enables frame to frame tracking of groups despite changes in global topology (here splitting).




Figure 1.9: Detecting and tracking group merging behavior. The level set formulation used here enables frame to frame tracking of groups despite changes in global topology (here merging).




Figure 1.10: Detecting and tracking group collision and splitting behavior. The level set formulation used here enables frame to frame tracking of groups despite changes in global topology (here merging and then splitting).

1.2.2 Hankel Operator based Anomaly Detection

For group anomaly detection in scenarios of crowded scenes, we have developed a macroscopic approach based on system identification techniques [c1]. In this approach we rely on low-level motion features, such as optical flow to extract relevant low order group dynamics from the video and use that to identify changes in group behavior. We assume that the low order dynamics of these motion features is governed by an unknown underlying Linear




Figure 1.11: Subplots a-b show optical flow for representative frames before and after panic behavior in UMN dataset. Subplot c shows the time evolution of how many normalized singular values (indicative of system order) exceed the 0.1 threshold, to determine the order of underlying LTI system. As the panic behavior develops, the system order drops drastically signaling the onset of an anomaly.

Time Invariant (LTI) system. To avoid computational difficulties in learning the underlying LTI model, we use subspace system identification techniques based on the Hankel matrix which can be constructed directly from the feature data. The spectral properties of Hankel matrix, and its range space encode useful information about the dynamics which can then be used to detect anomalous behavior [15, 6]. Specifically, we use a change in rank of the Hankel matrix to identify changes in the behavior. The matrix rank can be computed by singular value decomposition. In our application, the low level features typically lie in a high dimensional space (see example below); as a result the Hankel matrix can become very large, and SVD of large matrices can become a computational bottleneck. To expedite the SVD of the Hankel matrix, we use recently developed randomized algorithms [10] which use random sampling to identify a subspace that captures most of the action of a matrix, and efficiently obtain a low rank matrix approximations such as truncated SVD. Furthermore, we exploit the Teoplitz structure of the Hankel matrix while applying this randomized SVD approach, rendering the computation also memory efficient.

We have applied this methodology to robustly detect crowd panic behavior in the University of Minnesota (UMN) dataset [1]. We employ optical flow computed via the Lucas-Kanade [16] algorithm as the low level features. The optical flow is subsampled on a grid of 80×80 to capture coarse group behavior. Figure 1.11a-b shows the optical flow for two representative frames, one before and the other after the panic behavior emerges. As people move randomly before panic, the time evolution of optical flow is random. After the panic starts, the optical flow become more organized as people start exhibiting more directed motion. As a result, the LTI system representing the dynamics of optical flow would be high dimensional before the panic, and its order should drop as more organized behavior arises. We use a sliding window of $T = 100$ frames to construct the Hankel matrix, whose size therefore becomes $160,000 \times 50$. To determine the order of the underlying LTI system, we check how many the normalized singular values of Hankel matrix (computed using randomized SVD) exceed a prescribed threshold. Figure 1.11c shows the evolution of the order of

the system using a threshold of 0.1. Clearly, as the panic behavior arises the system order drops drastically signaling the onset of an anomaly.


1.2.3 Nonlinear Dynamical System Analysis for Dense Crowd Segmentation

In this section we consider the problem of segmenting highly dense crowded scenes which arise in public gatherings, such as one shown in the Figure 1.12a. Segmenting the scene into regions with distinct group motions/behaviors and characterization of this motion (e.g. formation of congestion, bottlenecks, etc), could enhance situational awareness of an analyst to preempt undesirable events such as public disturbance.


We have developed a nonlinear dynamical systems approach for robust dense crowd segmentation and characterization of internal dynamics within the segments [t2,p1]. A Hankel operator based analysis (as discussed in Section 1.2.2) reveals that a highly correlated motion exists in such dense crowd flows [c1]. Pushing this further, new insights can be obtained by treating dense crowds as a fluid flow driven by optical flow in the images. Given this analogy, one can then employ nonlinear dynamical systems techniques to detect coherent motion patterns in such flow fields and use them for crowd motion segmentation and change detection in crowd behavior.

The first step in such a nonlinear dynamical system analysis is to construct the flow map by advecting particles under the optical flow field. The flow map is then used for geometric, statistical, and spectral characterization of the crowd behavior. As shown in Figure 1.12b, a Finite Time Lyapunov Exponent (FTLE) approach detects coherent structure boundaries by computing extrema of maximum eigenvalues of the Cauchy Green deformation tensor [11]. On the other hand, eigenfunctions of the Perron Frobenius operator can be used to detect Almost Invariant Sets (AIS) which are regions with minimal leakage of trajectories in a statistical sense [8]. Figure 1.12c shows the AIS computed based on an Ulam approximation of the Perron Frobenius operator. Finally, Figure 1.12d shows the ergodic partitions (EP) obtained based on the eigenfunctions of the Koopman operator corresponding to the unit eigenvalue. To compute the ergodic partitions, we first construct the ergodic quotient by using time averages of spatial Fourier functions along particle trajectories, and then construct the diffusion coordinates based on Sobolev space norm of the negative index defined on the ergodic quotient [3]. Overall, each of these methods provide qualitatively similar segmentation of a crowd. However, note that the ergodic partitions also provide additional information on the internal structure of the flow: for example the areas where congested flow transitions to a more free flowing crowd is highlighted by yellow/blue colors.


We have also explored the use of Koopman Mode Analysis (KMA) [4, 18] for a model free spectral characterization of group behavior. Here changes in Koopman spectra over time can be used to signal when the change occurs in behavior, while the changes in the spatial Koopman modes highlight regions in the image where this change most likely occurred [t2].




(a) Mecca dense crowd sequence



(b) Finite Time Lyapunov Exponent



(c) Almost Invariant Set



(d) Ergodic Partitions

Figure 1.12: Application of different nonlinear dynamical system techniques for crowd segmentation. b) shows FTLE field with high values indicating boundaries of coherent structure. c) shows the AIS highlighted with different colours. d) shows the EP. While FTLE and AIS identify regions of different qualitative crowd motion, EPs in addition also provide information on the internal structure of the flow: for example the areas where congested flow transitions to a more free flowing crowd is highlighted by yellow/blue colors in subplot d.

1.2.4 Koopman Operator based Nonlinear Dynamic Textures

In this section we describe another application of Koopman Mode Analysis (KMA) for modeling more general video scenes with statistically repetitive spatiotemporal patterns which are referred to as Dynamic Texture (DTs) in the computer vision literature [c5].


A popular generative modeling paradigm for DTs is to treat it as a sample output of a stochastic linear dynamical system (LDS). Despite their simplicity, such Linear Dynamic Texture (LDT) models have shown to be surprisingly useful in domains such as video synthesis, classification, and segmentation. Experimental evidence however shows that LDTs are sometimes inadequate to effectively describe the time evolution of the real world dynamic scenes which exhibit globally nonlinear dynamics; nonlinear correlation between frames due to complex motion, such as chaotic motion or camera motion; sudden changes in scene due to depths discontinuities, occlusions, etc.; and coexisting multiple regions belonging to a semantically different visual process. To address these limitations, many LDT variants have been proposed in the literature (see [c5] for details), but none of these methods give full nonlinear treatment of DTs.

We have developed a nonlinear dynamic texture model in which both the state transition and observation function are nonlinear [c5]. Our approach is based on KMA which uses Koopman spectral decomposition to determine a data-driven modal decomposition and model reduction [18]. The Koopman operator is a linear but infinite-dimensional operator whose modes and eigenvalues capture the evolution of observables (e.g. video frames in the DT case) describing any underlying (nonlinear) dynamical system. We exploit this aspect in constructing a linear stochastic system in Koopman mode space and propose it as a generative model for nonlinear DTs. We refer to this model as Koopman Mode Dynamic Texture (KMDT). We use a sparse Dynamic Mode Decomposition [12] based numerical procedure for KMA to learn the KMDT model from videos.


We compared KMDT with LDT and kernel dynamic texture (KDT) approaches on several complex real world videos, and find superior modeling performance (see Figure 1.13).



(a) Video 1



(b) Video 1




(c) Video 1



(d) Video 3



(e) Video 3



(f) Video 3

Figure 1.13: Snapshots from different texture videos and comparison of modeling accuracy of KMDT with LDT and KDT approaches as a function of different number of modes retained. The accuracy is measured in terms of average PSNR, see [c5] for details. KMDT provides better accuracy compared to other methods for same number of modes retained.

Chapter 2

Transitions at UTRC

2.1 Data Driven Nonlinear Model Reduction for PHM/Big Data Applications

The Koopman Mode Analysis based nonlinear model identification/reduction concept developed for dynamic texture modeling (as described in Section 1.2.4) was applied in an internally funded project to learn dynamic models for load estimation in rotorcraft prognostic and health management (PHM) applications. Currently, we are also exploring other applications of this approach including optimal sensor selection and big data streaming analytics. POC Andrzej Banaszuk, UTRC, 860-610-7381.

Chapter 3

Personnel Supported

UTRC personnel: Amit Surana and Kunal Srivastava.

Chapter 4

Publications

Journal Papers in Preparation

[p1] A. Surana, and I. Mezić, “Detecting Coherent Structures in Crowd Videos”, to be submitted to Physica D.

[p2] M. Niethammer, A. Surana and A. Tannenbaum, “Detecting and Tracking Groups in Crowd Videos”, in preparation.

Conference Papers

[c1] A. Surana, A. Nakhmani and A. Tannenbaum, “Dynamical Systems Framework for Anomaly Detection in Videos”, Conference on Decision and Control, 2013.

[c2] A. Surana “Unsupervised Inverse Reinforcement Learning with Noisy Data”, Conference on Decision and Control, 2014.

[c3] A. Surana, and K. Srivastava, “Bayesian Nonparametric Inverse Reinforcement Learning for Switched Markov Decision Processes”, International Conf. on Machine Learning and Applications, 2014.

[c4] A. Nakhmani, A. Surana, and A. Tannenbaum, “Macroscopic Analysis of Crowd Motion in Video Sequences”, Conference on Decision and Control, 2014.

[c5] A. Surana “Koopman Operator Based Nonlinear Dynamic Textures”, submitted ACC, 2015.

Invited Sessions

The following invited session was organized with AFOSR support and contain AFOSR-funded papers:

-2013 SIAM Conference on Control and Applications, San Diego: Dynamical System and Control Based Methods for Computer Vision Problems; Organizers: A. Surana and A. Tannenbaum.

Talks

[t1] A. Surana “Dynamical Systems Framework for Anomaly Detection in Videos”, presented in SIAM Conference on Control and Applications, San Diego, July 2013.

[t2] A. Surana, “Dynamical System Analysis of Crowd Videos”, presented in BIRS Workshop on Uncovering Transport Barriers in Geophysical Flows, Banff, Sep 2013.

Bibliography

- [1] Umn database (detection of unusual crowd activity). <http://mha.cs.umn.edu>.
- [2] Pieter Abbeel and Andrew Y. Ng. Apprenticeship learning via inverse reinforcement learning. In *ICML*, 2004.
- [3] Marko Budisic and Igor Mezic. Geometry of the ergodic quotient reveals coherent structures in flows. *Physica D: Nonlinear Phenomena*, 241(15):1255–1269, 2012.
- [4] Marko Budisic, Ryan M. Mohr, and Igor Mezic. Applied koopmanism. *arXiv:1206.3164v1*, 2012.
- [5] Jaedeug Choi and Kee-Eung Kim. MAP Inference for Bayesian Inverse Reinforcement Learning. In *NIPS*, pages 1989–1997, 2011.
- [6] Tao Ding, Mario Sznajder, and Octavia I. Camps. Fast track matching and event detection. In *CVPR*, 2008.
- [7] E. B. Fox. *Bayesian Nonparametric Learning of Complex Dynamical Phenomena*. Ph.D. thesis, MIT, Cambridge, MA, 2009.
- [8] G. Froyland and K. Padberg. Almot-invariant sets and invariant manifolds-connecting probabilistic and geometric descriptions of coherent structures in flows. *Physica D*, 238.
- [9] Samuel J. Gershman and David M. Blei. A tutorial on bayesian nonparametric models. *Journal of Mathematical Psychology*, 56:1–12, 2012.
- [10] Nathan Halko, Per-Gunnar Martinsson, and Joel A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53(2):217–288, 2011.
- [11] G. Haller. Distinguished material surfaces and coherent structures in 3d fluid flows. *Physica D*, 149.
- [12] M. R. Jovanovic, P. J. Schmid, and J. W. Nichols. Sparsity-promoting dynamic mode decomposition. In *arXiv:1309.4165v1 [physics.u-dyn]*, 2013.
- [13] Kris M. Kitani, Brian D. Ziebart, James Andrew Bagnell, and Martial Hebert. Activity Forecasting. In *ECCV (4)*, pages 201–214, 2012.

- [14] T. L. Lai. Sequential analysis: Some classical problems and new challenges. *Statistica Sinica*, 11:303–408, 2001.
- [15] Binlong Li, Mustafa Ayazoglu, Teresa Mao, Octavia I. Camps, and Mario Sznajder. Activity recognition using dynamic subspace angles. In *CVPR*, pages 3193–3200, 2011.
- [16] B. D. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. *Proceedings of Imaging Understanding Workshop*, pages 121–130, 1981.
- [17] S. P. Meyn, A. Surana, Y. Lin, and S. Narayanan. Anomaly detection using mori-zwanzig-shannon Markov models in a distributed sensor network. *CDC*, 2009.
- [18] Igor Mezic. Analysis of fluid flows via spectral properties of the koopman operator. *Annual Review of Fluid Mechanics*, 45, 2012.
- [19] O. Michailovich, Y. Rathi, and A. Tannenbaum. Image segmentation using active contours driven by the bhattacharyya gradient flow. *IEEE Trans. Image Processing*, 16:2787–2801, 2007.
- [20] R. M. Neal. Markov Chain Sampling Methods for Dirichlet Process Mixture Models. *Journal of Computational and Graphical Statistics*, 9(2):249–265, 200.
- [21] Gergely Neu and Csaba Szepesvári. Apprenticeship Learning using Inverse Reinforcement Learning and Gradient Methods. In *UAI*, pages 295–302, 2007.
- [22] Gergely Neu and Csaba Szepesvári. Time parsers by inverse reinforcement learning. 77(2), 2009.
- [23] Andrew Y. Ng and Stuart J. Russell. Algorithms for Inverse Reinforcement Learning. In *ICML*, pages 663–670, 2000.
- [24] Marc Niethammer, Patricio A. Vela, and Allen Tannenbaum. Geometric observers for dynamically evolving curves. volume 30, 2008.
- [25] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley and Sons, Inc. Hoboken, N. J., 1994.
- [26] Deepak Ramachandran and Eyal Amir. Bayesian Inverse Reinforcement Learning. In *IJCAI*, pages 2586–2591, 2007.
- [27] Y. Rathi, N. Vaswani, and A. Tannenbaum. A generic framework for tracking using particle filtering with dynamic shape prior. *IEEE Trans. Image Processing*, 16:1370–1382, 2007.
- [28] Y. Rathi, N. Vaswani, A. Yezzi, and A. Tannenbaum. Tracking deforming objects using particle filtering for geometric active contours. *IEEE Pattern Analysis and Machine Intelligence*, 29:1470–1475, 2007.

- [29] N. D. Ratliff, J. A. Bagnell, and M. A. Zinkevich. Maximum margin planning. In *ICML*, 2006.
- [30] A. Surana, T. A. Frewen, A. Banaszuk, H. Sane, and S. Narayanan. Motion planning of uav assets to identify threats in uncertain urban environment. *Special Operations Forces Industry Conference (SOFIC)*, 2009.
- [31] Richard S. Sutton, Doina Precup, and Satinder P. Singh. Between MDPs and Semi-MDPs: A Framework for Temporal Abstraction in Reinforcement Learning. *Artif. Intell.*, 112(1-2):181–211, 1999.
- [32] Umar Syed, Michael H. Bowling, and Robert E. Schapire. Apprenticeship learning using linear programming. In *ICML*, pages 1032–1039, 2008.
- [33] N. Vaswani, Y. Rathi, A. Yezzi, and A. Tannenbaum. Deform pf-mt (particle filter with mode tracker) for tracking nonaffine contour deformations. *IEEE Trans. Image Processing*, 19:841–857, 2010.
- [34] Namrata Vaswani. Additive change detection in nonlinear systems with unknown change parameters. *IEEE Transactions on Signal Processing*, 55(3):859–872, 2007.
- [35] Brian D. Ziebart, Andrew L. Maas, J. Andrew Bagnell, and Anind K. Dey. Maximum Entropy Inverse Reinforcement Learning. In *AAAI*, pages 1433–1438, 2008.