

Intention-based Behavioral Anomaly Detection

Fan Hung^{1*} Xu Xie^{1*} Andrew Fuchs^{2*}

Michael Walton² Siyuan Qi¹ Yixin Zhu¹ Doug Lange² Song-Chun Zhu¹ * † ‡

Abstract

In this work we consider the problem of detecting anomalous deviations from expected behavioral trajectories generated by a multi-agent system. Common methods for time series anomaly detection do not explicitly consider the goal-directed nature of rational agents. This work proposes a method for detection of anomalous behaviors based on agent intent formulated using agent-based Lagrangian Mechanics. We propose an anomaly detection method that simultaneously learns to 1) predict the intended goals of agents from their trajectories, and 2) detect anomalies based on the predictions. The method uses a structured one-class support vector machine (structured OC-SVM), where latent variables represent the intentions (goals) of agents, and SVM weights represent the attraction or repulsion strength of potential goals and obstructions. Model parameters are learned in an unsupervised way. We conduct experiments in a marine surveillance setting, where we monitor high-traffic ports and detect anomalous vessels. Experimental results show that our algorithm shows promise for detecting anomalous behaviors in systems of goal-directed agents.

Introduction

We considered the coupled problems of intention prediction and anomaly detection on trajectories. Anomaly Detection is the problem of finding patterns in data that differ from the majority of normative support. This work is primarily motivated by an application to maritime vessel traffic surveillance. The United State Coast Guard (USCG) must monitor and patrol numerous high traffic ports and waterways for indications of unusual, suspicious or potentially malicious behavior. The human-labor intensive process of discriminating anomalous vessel traffic from normative background often relies on years of training and intuition. While the anomalous vessel trajectories of interest occur very infrequently, voluminous historical data of normal traffic is publicly available. Therefore, we pose the problem of detecting suspicious

** Fan Hung, Xu Xie and Andrew Fuchs contributed equally to this work

†¹ UCLA Center for Vision, Cognition, Learning, and Autonomy (VCLA) at Statistics Department.

‡² Space and Naval Warfare Systems Center Pacific (SSC PAC). Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

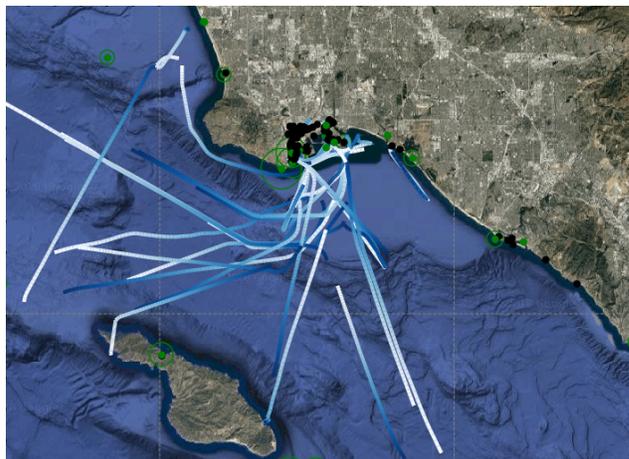


Figure 1: AIS Data Trajectories. Color Darkening in trajectories indicate forward passage of time.

vessels as anomaly detection given a data set of assumed normative traffic.

Approaches for trajectory anomaly detection typically construct a model of the data distribution for example, using DBSCAN, Gaussian mixture models, nearest neighbors, or One Class Support Vector Machines (OC-SVM). Operating directly on a trajectory treated as a vector representation of a sequence of points encounters problems such as high dimensionality for longer sequences and is not robust to noise or time shift. Many works construct more compact representations of trajectories but rely on explicit spatio-temporal modelling. These approaches have several shortcomings. They do not capture higher level semantics and place greater burden on human operators to interpret detected anomalies before raising alerts.

A naturally related field is that of Intention Prediction, which is the task of explaining observed human behavior and determining motivations. The challenge and draw of this vein of work is that human intention is difficult to define and often requires contextual knowledge beyond what is typically present in numerical data. In line with such work, we aim to reason about the agent planning processes behind trajectories to infer the intentions of observed agents, and then detect intent-based anomalies. Explicitly modelling

path planning exploits the underlying semantic structure of agent trajectories and presents detections which are intuitively interpretable.

To do so, our model combines a formulation of path intentions which detects goals for pedestrians walking through a scene with the Structured OC-SVM following the example of the Hidden Markov Model One Class Support Vector Machine (HMM-OC-SVM) (Görnitz, Braun, and Kloft 2015). The HMM-OC-SVM is an anomaly detector which learns the transition matrix, emission probabilities, and prior distribution for sequential data generated by an HMM. It couples the HMM Viterbi optimization with OC-SVM training optimizations, and at training time alternates the two optimizations to infer latent state transition chains and then infers a decision boundary which measures the degree of conformity to the HMM. Similarly, we couple and alternate optimization goal choice for trajectories and optimization of a SVM decision boundary that measures conformity.

Our intention prediction model follows agent-based Lagrangian Mechanics, which is related to potential-based path planning (Xie et al. 2018). This models trajectories as responding to attractive and repulsive fields of point sources similar to those of stationary charged particles. We allow the magnitudes of the attractive forces of the particles to vary and construct feature vectors that such that OC-SVM training fits these magnitudes and determines popular goals from a set of candidates.

We test our algorithm in a marine surveillance setting using marine Automatic Identification System (AIS) data recorded near ports. Experiment results show that our algorithm outperforms standard kernel-based support vector machines, and that an intent-based formulation has potential in this domain.

Related Work

Intention Prediction This is a relatively new but important domain in artificial intelligence with many potential applications, and most work is focused on predicting human intentions (Yuen and Torralba 2010; Wang et al. 2012; Li and Fu 2014; Vu et al. 2014; Qi et al. 2017; Rhinehart and Kitani 2017; Qi, Jia, and Zhu 2018), Particularly relating to trajectory data, several works model agent path planning-based on a valuation of their surroundings. Inverse reinforcement learning is also used to infer a valuation of walkable space (Ziebart et al. 2009; Kitani et al. 2012). One formulation constructs a state space from segmented images, generates a state valuation based on frequency of traversal, and predicts paths that traverse the space (Walker, Gupta, and Hebert 2014). Our work builds off of a formulation of trajectory intention which explicitly models path planning for explicit goals and boundaries using vector fields (Xie et al. 2018). Other works model the interaction between pedestrians in crowded scenes (Kuderer et al. 2012; Alahi et al. 2016; Ma et al. 2017). We also include similar interactions between agents.

Anomaly Detection General data driven Anomaly Detection considers separating vector spaces into inlier and out-

lier regions or direct data mining (Chandola, Banerjee, and Kumar 2009; Zimek, Schubert, and Kriegel 2012; Agrawal and Agrawal 2015). These works focus on aspects such as varied density of inlier clusters, high dimensionality, decision boundary formalization, graphical structure, and mutual information. Anomaly Detection also receives attention in discrete domains for the purposes of network intrusion and fraud detection (Chandola, Banerjee, and Kumar 2012; Akoglu, Tong, and Koutra 2015; Buczak and Guven 2016). Many of these works study sequential data, but the discrete formulation does not match our setting. Our approach similar to those which construct SVM kernels to define distance between discrete sequences. Recent works in anomaly detection have also considered spatio-temporal anomaly detection while modelling human behavior in the domain of security camera footage (Li, Mahadevan, and Vasconcelos 2014). Many works focus on detecting anomalous video frames using imaging techniques. Some of these approaches model the behavior of pedestrians and capture intention prediction by including path planning in formulations.

In maritime anomaly detection, many approaches focus on geographic modelling such as mining frequently travelled routes and path segments. (Sidib and Shu 2017). Route-based approaches for the most part declare anomalous behaviour based on distance from routes. There is a previous, similar potential field-based model, but the field is constructed differently and is not goal-based (Osekowska and Carlsson 2015). Some other works build rule-based systems for detecting anomalies, but many of these do not model path planning.

Our work builds off a stream of structured prediction-based approaches using variations of one-class support vector machines (OC-SVM) (Muller et al. 2001). Using kernel-based techniques, these methods learn a representation of the data set in a transformed feature space. As an extension to methods that assume the data are independently distributed, some approaches exploit additional latent structure of the data. Some for example, uncover the underlying grammar structure of speech or text for intrusion detection (Joachims et al. 2009; Rieck et al. 2010). We follow the example of a hidden Markov kernel for anomaly detection in data with latent temporal dependency structure (Görnitz, Braun, and Kloft 2015).

SVM Anomaly Detection

In this section, we introduce the formulation of the standard one-class support vector machine (OC-SVM), which will later be extended to latent structure-based OC-SVM. The OC-SVM finds a linear separation between vector representations of normal data and anomalies. The standard OC-SVM assumes that the data is independently distributed, while the latent OC-SVM first uncovers underlying latent variables and then performs anomaly detection.

The OC-SVM derives from a similar formalization to the standard SVM classifier. In short, the standard two-class SVM computes a linear boundary by maximizing the boundary's distance to a subset of the data points closest to the true border between the two classes. It optimizes following the expectation maximization method by iterative fit-

ting of boundary parameters and re-selection of the subset of closest points. Practically, the two classes' data overlap, so the optimization problem is relaxed and a few training mis-classifications are allowed. The subset of closest data points is the set of "support vectors" and lies on both sides of the boundary.

The OC-SVM forms an outlier detector from this framework. Its boundary tightly encompasses the given data, dividing the two classes: "outlier" and "inlier". For a given data set, it chooses a fraction of the support vectors (parametrized by ν) to lie on the outlier side of the boundary. These are typically trained such that the outliers end up on the same side of the boundary as the origin. This can be interpreted as a threshold on a scoring function which indirectly measures the likelihood that a new data point belongs to the same distribution. More formally, for data vectors $x_1, \dots, x_N \in X \subseteq \mathbf{R}^D$ embedded in a feature space using mapping $\phi(x_i) : \mathbf{R}^D \rightarrow \mathbf{R}^F$, the OC-SVM has the objective to linearly separate the data into normal data and anomalies:

$$\begin{aligned} \min_{\omega, \rho, \xi \geq 0} \quad & \|\omega\|^2 - \rho + \frac{1}{\nu N} \sum_{i=1}^N \xi_i \\ \text{subject to} \quad & \langle \omega, \phi(x_i) \rangle - \rho + \xi_i \geq 0, \forall i = 1, \dots, N \end{aligned} \quad (1)$$

where $\langle \cdot, \cdot \rangle$ denotes inner product, $\omega \in \mathbf{R}^F$ is the weight vector of the linear model, $\rho \in \mathbf{R}$ is the threshold, and $\xi_i \in \mathbf{R}$ are the slack variables allowing for class overlap. The boundary is a hyper-plane in the earlier linear case described, but here, using the mapping ϕ and the kernel method, the decision boundary of the SVM can be made nonlinear. With a radial basis kernel, for example, $\langle \phi(x), \phi(y) \rangle = K(x, y) = \exp(-\gamma \|x - y\|^2)$ a decision boundary arises that resembles that of a nearest neighbors approach.

Intent Prediction by Agent-based Lagrangian Mechanics

Lagrangian Mechanics

Lagrangian Mechanics reformulates Newtonian balancing of forces into an optimization problem. It models physical interactions as a minimization problem over possible paths a system can take and postulates that situations unfold following minimizing paths. For example, a ball rolling through a hilly landscape will roll through following a path which stays as much as possible within the valleys. Let such a particle's position over time t be $x(t)$, and velocity $\dot{x}(t)$, and the force field affecting it be $\vec{F}(x(t))$. The objective function of minimization is called action and is defined as the integral over time of the so called Lagrangian function. This problem can be written as:

$$\Gamma(t_1, t_2) = \operatorname{argmin}_x \int_{t_1}^{t_2} L(x, \dot{x}, t) dt \quad (2)$$

Where t_1 and t_2 are the start end of the time that we are modelling over, $\Gamma(t_1, t_2) \ni x(t) \forall t \in [t_1, t_2]$ is the particle's trajectory and $L(x, \dot{x}, t)$ is the Lagrangian function for

this scenario. The Lagrangian is the kinetic energy minus the potential energy of the particle:

$$L(x, \dot{x}, t) = \frac{1}{2} m \dot{x}(t)^2 - \int_x \vec{F}(x(t)) d\vec{x}(t) \quad (3)$$

Here, m is the particle's mass, and the integral is taken over the particle's path. The integral term can also be recognized as a work done against the potential field over the path of the particle. In our model this term will be parameterized and fitted.

Agent-based Lagrangian Mechanics

Agent-based Lagrangian Mechanics models agents as particles following force fields. These force fields are generated by attractive or repulsive point source goals or obstructions. In a maritime setting, particles correspond to vessels, goals to points of interest such as ports or fishing spots, and repulsive sources to obstacles such as landmasses or other vessels. In physics, attraction and repulsion forces are modelled as an inverse squared distance relationship. Symbolically, the magnitude of the force F of attraction between two objects at positions x_1 and x_2 can be modelled as:

$$F_{1,2} = k_f \frac{k_1 k_2}{\|x_1 - x_2\|_2^2} \quad (4)$$

where k_f, k_1, k_2 are scalar constants unique to the type of force, and the nature of the objects at positions 1 and 2 respectively. k_1 and k_2 describe the inherent attractive strength of the objects. They are mass when the force is gravitational or charge when the force is electrical. The direction of the force is given by the difference in positions x_1 and x_2 and may be flipped when the force is repulsive. Agents choose the fields they respond to, but there is no choice but to avoid any obstacles. In the maritime example, vessels move towards one point of interest at a time and need to avoid collisions with any vessels or landmasses. A discrete set of points of interest represents a higher level action space and set of available intents. Predicting an agent's goal and intent is choosing the source of attraction that minimizes the Lagrangian objective for a given trajectory. Goals are chosen one at a time, and all repulsive sources are shared by all agents. This is also extended to allow sequences of goals.

Intention-based Anomaly Detection

Intention Prediction

We denote a set of M agents as $A = \{a_i : i = 1, \dots, M\}$ and G sources of goals which agents may choose from. Agents plan paths around obstacles, so we include R constantly present repulsive sources representing the environmental obstructions or even other non-stationary agents. The total set of point sources is $S = \{s_j : j = 1, \dots, G + R\}$.

When $a_i \in A$ selects goal $s_j; j \in \{1, 2, \dots, G\}$ the net force which a_i is affected by can be expressed as

$$\vec{F}(x) = \vec{F}_{a_i, s_j}(x) + \sum_{k \in R} \vec{F}_{a_i, s_k}(x) \quad (5)$$

In the work of Xie et al, the force magnitudes are inversely proportional to the Mahalanobis distances between the agent

and the attracter/repeller (Xie et al. 2018). Hence, we have $\|\vec{F}_{a_i, s_j}(x)\| \propto ((x - s_j)^T \Sigma^{-1} (x - s_j))^{-1}$. For simplicity, we use a diagonal covariance matrix, $\Sigma = \sigma^2 \mathbf{I}$. The agents respond to the same goals identically, or in the maritime setting, all vessels have the same attraction toward a given port, but ports may vary between each other in attractiveness. We model the action of constants k_f and k_1 with a constant ω_j and the action of constant k_2 with a constant σ . Each goal's ω_j will be learned and σ will be a hyper-parameter. Repulsion forces act in the same way, but their vector directions are flipped, or equivalently, the corresponding ω_j will be negative. We rewrite the force:

$$\|\vec{F}_{a_i, s_j}(x)\| = \omega_j \|\vec{F}_j(x)\| = \omega_j \sigma^2 \|x - s_j\|^{-2}. \quad (6)$$

Two different σ 's are used: one each for the set of obstacles and the set of attracters. The σ for repulsion is smaller as agents are driven toward goals from far away but may only respond to obstacles close by when there is a chance of collision. When agent a_i is assumed to choose only one goal through its trajectory, we can model its choice as follows. a_i 's trajectory in continuous space, moving toward an attractive goal s_j , over discrete time steps is denoted as a sequence of points $\Gamma_{i,j}(t_1, t_2) = (x(t_1), \dots, x(t_2))$. To predict an agent's intent, we choose the goal j whose associated action is minimal and best fits the trajectory $\Gamma_{i,j}$. At the large geographic scale we test in, we assume vessels change speed negligibly, so that the velocity, \dot{x} is relatively constant, optimization of the Lagrangian objective becomes a problem of minimizing the work done against the force field. In a discrete time setting, we can rewrite the optimization as:

$$\operatorname{argmax}_j \sum_t \langle \omega_j \vec{F}_j(x), \Delta \vec{x}(t) \rangle \quad (7)$$

with $\Delta \vec{x}_t = \vec{x}_{t+1} - \vec{x}_t$. This has the probabilistic interpretation that the likelihood of a given agent trajectory is:

$$P(\Gamma(t_1, t_2)) \propto \frac{1}{Z} \exp\left(\sum_{t=t_1}^{t_2-1} \langle \omega_j \vec{F}_j(x), \Delta \vec{x}(t) \rangle\right) \quad (8)$$

The goal is chosen to maximize the log probability of the trajectory given that the agent is following the force field towards it, $P(\Gamma(t_1, t_2)|s_j)$. Agents may also move through a sequence of goals or change goals. For example, a ferry may move back and forth between two ports, so we include, in the model, sequences of goals, $s_j(t)$. We model the probability that an agent continues heading toward its current goal through a timestep as κ and the probability that it changes to another goal with even chance over other goals totaling to $1 - \kappa$. More explicitly,

$$P(s(t)|s(t-1)) = \begin{cases} \frac{1-\kappa}{N-1}, & s(t) \neq s(t-1) \\ \kappa, & s(t) = s(t-1) \end{cases} \quad (9)$$

To incorporate this, first, consider the trajectory likelihood as a product over timesteps:

$$\frac{1}{Z} \prod_{t=t_1}^{t_2-1} \exp(\langle \omega_j \vec{F}_j(x), \Delta \vec{x}(t) \rangle) = \frac{1}{Z} \prod_{t=t_1}^{t_2-1} P(\Gamma(t, t+1)|s_j) \quad (10)$$

We can then write $P(\Gamma(t_1, t_2)|s(t))$:

$$\propto \frac{1}{Z} \prod_{t=t_1}^{t_2-1} P(\Gamma(t, t+1)|s(t)) P(s(t)|s(t-1)) \quad (11)$$

where we define $P(s(t_1)|s(t_1-1)) := P(s(t_1))$ as the prior on the initial goal. This makes the new objective:

$$\operatorname{argmax}_{j(t)} \sum_t \langle \omega_{s(t)} \vec{F}_{s(t)}(x), \Delta \vec{x}(t) \rangle + \log P(s(t)|s(t-1)) \quad (12)$$

This optimization is a choice of a series of goals, which can be computed with dynamic programming. The choice of goals is extended to allow for a sequence of goals with a regularization on the number of switches. This way, in the maritime setting, we can model, for example, a ferry moving between two ports. In a sense, the agent's sequence of goal switches behaves like a hidden markov model with a known transition matrix that has κ along the diagonal and $\frac{1-\kappa}{N-1}$ in all other entries. From this perspective, the incremental probabilities, $P(\Gamma(t, t+1)|s(t))$ are akin to emission probabilities. This formulation differs slightly from Xie et al's which models a probability over number of switches and imposes a bound on the number of switches (Xie et al. 2018). This is like the regime changes from switching state space models (Nicholas Hoernle 2018). In our work, goals define regimes, but may be revisited. The main difference from that of Hoernle et al is that our "emission" probabilities do not directly use a state space model. Also, this approach also does not label any time steps in a semi-supervised fashion.

Latent OC-SVM Objective

For SVM problems that exhibit underlying latent structures, it is desirable to first infer hidden variables, based on which we can do further analysis. For example, in the maritime setting, agents' goal choices make up their latent intents. In the anomaly detection literature, an extension to the OC-SVM is the latent OC-SVM (Görnitz, Braun, and Kloft 2015). The data vectors in this treatment belong to an observable subspace. The remainder of the space is inferred by a separate maximization of likelihood of the latent structure. This approach uses a joint feature map $\Psi(x, z)$ in place of $\phi(x)$ in the original OC-SVM, which allows for a latent variable z to be assigned to each data point. The observable subspace will derive from the force field-based probability of the trajectory, and indicator variables will be used to represent the inferred goal choices. The latent variable assignment is determined by $\operatorname{argmax}_z \langle w, \Psi(x, z) \rangle$. Formally, the objective of the latent OC-SVM is:

$$\begin{aligned} \min_{\omega, \rho, \xi \geq 0} \quad & \|\omega\|^2 - \rho + \frac{1}{\nu n} \sum_{i=1}^M l(\xi_i) \\ \text{subject to} \quad & \max_{z \in \mathcal{Z}} (\langle \omega, \psi(x_i, z) \rangle + \delta(z)) - \rho + \xi_i \geq 0, \\ & \forall i = 1, \dots, M \end{aligned} \quad (13)$$

where l is a loss function, which is often defined as the monotonically non-decreasing hinge loss function $l : l(t) = \max(0, t)$.

Often, this optimization is interpreted as maximizing a log probability of the observation and the hidden variables, *i.e.*, $\log(p(x|z)) = \langle w, \Psi(x, z) \rangle$, and $\log(p(z)) = \delta(z)$. For example, Görnitz et al uses x as a sequence of emissions of a hidden markov model, z as a latent state transition chain, and w as an embodiment of transition and emission probabilities (Görnitz, Braun, and Kloft 2015). There, latent state transition chains were chosen using the viterbi algorithm. In general, to train such a model, OC-SVM training and optimization of latent vector choices $\operatorname{argmax}_z \langle w, \Psi(x, z) \rangle$ are alternated. We will use this extension, *i.e.*, the latent OC-SVM, to add agent’s choice of goals. Specifically, the latent variable z represents the intention/goal of an agent, and the intention prediction can be formulated as:

$$\begin{aligned} z &= \operatorname{argmax}_{z \in \mathcal{Z}} \log p(z|x) \\ &= \operatorname{argmax}_{z \in \mathcal{Z}} \log p(x|z) + \log p(z) \\ &= \operatorname{argmax}_{z \in \mathcal{Z}} \langle \omega, \Psi(x, z) \rangle + \delta(z) \end{aligned} \quad (14)$$

We calculate our feature $\Psi(x, z)$ using indicator variables to represent goal choices as follows: Notationally, let $\mathbf{1}[F]$ be the indicator of F equalling 1 if the condition represented by F is satisfied and 0 otherwise. Let $\mathbf{1}[z_t = j]$ be 1 if the agent’s current chosen goal is s_j . Let $\mathbf{1}[z_{t+1} = l \wedge z_t = j]$ be 1 if the agent switches from goal s_j to goal s_l between times t and $t + 1$. And let $\mathbf{1}[z \in B]$ be 1 if z indexes an obstacle. We can write the objective as:

$$\begin{aligned} &\sum_{t=2}^{T-1} \sum_{k \in [N+B]} (\mathbf{1}[k = j] + \mathbf{1}[k \in B]) * \omega_k \langle \vec{F}_k(x_t), \Delta \vec{x}_t \rangle \\ &+ \sum_{t=3}^{T-1} \sum_{k, l \in [N], k \neq l} \mathbf{1}[z_{t+1} = l \wedge z_t = j] \log \left(\frac{1 - \kappa}{N - 1} \right) \\ &+ \sum_{t=3}^{T-1} \sum_{k \in [N]} \mathbf{1}[z_{t+1} = k \wedge z_t = k] \log(\kappa) \end{aligned} \quad (15)$$

The first sum can be taken to represent $\langle \omega, \Psi(x, z) \rangle$. We can rewrite it as:

$$\sum_{k \in [N]} \omega_k * \mathbf{1}[k = j] * \mathbf{1}[k \in N^-] * \sum_{t=2}^{T-1} \langle \vec{F}_k(x_t), \Delta \vec{x}_t \rangle \quad (16)$$

Here, the k -th term of the sum gives k -th component of ω and Ψ . Each component represents the total work done by the attractive field of a goal. The remaining two sums representing goal switching can be taken to represent a regularization term. κ close to 1 penalizes switching too often.

In the learning stage, we optimize the parameters in an iterative way. We first fix the force magnitudes ω and solve the best z for each trajectory using dynamic programming. Then we solve the ω using OC-SVM. We repeat the above process until it converges in a similar fashion to that for the latent OC-SVM (Görnitz, Braun, and Kloft 2015).

Experiments

Marine Surveillance

In this study we utilized historical Automatic Identification System (AIS) data provided by marinecadastre.gov which includes position, course and speed information for a variety of marine vessel types (Bureau of Ocean Energy Management (BOEM) and National Oceanic and Atmospheric Administration (NOAA)). The data was reported in years 2011 and 2013 and was not limited to a specific vessel type. To reduce computational complexity and increase the likelihood of higher density data, we limited our scope to AIS data gathered in higher traffic marine regions surrounding San Diego, CA and Long Beach, CA. Additionally, we incorporated data from osav-usdot.opendata.arcgis.com for known port locations, which we narrowed by the bounding regions for the two test locations (U.S. Department of Transportation (USDOT) / Bureau of Transportation Statistics (BTS)).

Synthetic Anomalous Data

In this study, we train an unsupervised model, assuming no access to anomalous ships at training time; to support quantitative evaluation of our model, we generated synthetic anomalous trajectories based on adding realistic deviation to historical AIS traffic and modifications from domain knowledge provided by maritime surveillance subject matter experts.

We generated test data using statistics gathered from the historical AIS data. The historical data from bounding boxes around the two test settings were extracted and then split into independent tracks for each vessel. In addition to the location information, vessel type, vessel type frequency, and vessel speed over ground (SOG) data was recorded for each AIS entry. Traversal frequencies were calculated over discretized grids of the settings. Land areas were labelled using coastline data. Also, start and end points for trajectories were recorded to represent high interest regions.

To generate a trajectory, we utilize standard A^* shortest path search on randomly selected start and end points and randomly perturb the path. The start and end point are chosen randomly from the grid with probability proportional to the traversal frequency. For a portion of our synthesized trajectories, these points are restricted to grid squares along coastlines and chosen with the previously mentioned probabilities. Next, an A^* search was run using the chosen start and end. During search, the cost of grid spaces was randomly modified. With a set probability at each step, the traversal frequencies multiplied by a scale factor are added to the incremental cost of generated child nodes. The probability that this frequency penalty is added and the discount scale factor are left as parameters. This penalty enforces that paths behave less like previous history. In addition to the discount factor and inclusion probability, path modifiers were included allowing for sudden divergence from a previously expected trajectory to a new destination, path smoothing to reduce jaggedness, or loitering. The modifiers provide additional sources of anomalous behavior, which serve to model behaviors of interest listed by USCG, namely unexplained

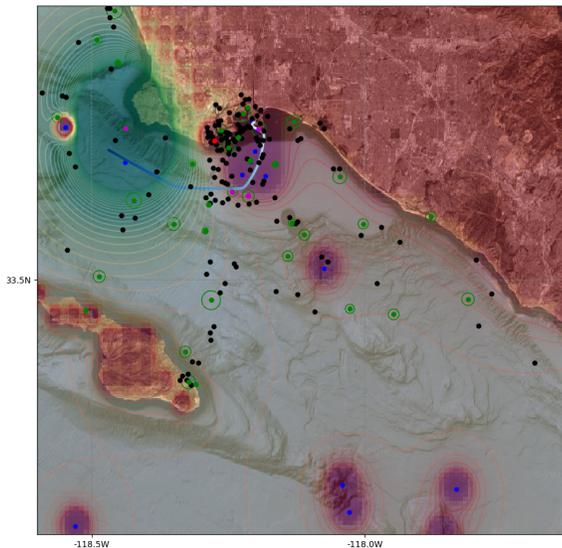


Figure 2: Long Beach Heatmap and One Prediction. Attraction heat increasing from red to green. The goals corresponding to lower 10-th percentile of attraction weights in black, repelling goals with negative weight red, and attractive goals in green and plotted with circle with radius proportional to attraction weight. Background obstacle ships in blue. Trajectory of interest darkening with passage of time. Goals predicted for trajectory in purple

loitering and unusual ports. For initial testing and experimentation, the modifiers were limited to a single instance for a given path. The loitering modifier added multiple position reports randomly generated in the selected end point region, with the region selected based on the frequency data. Additionally, the loitering points are added to the ends of complete paths. The sudden divergence modifier truncates the original path’s end section and appends a new path from the truncation point to the divergent endpoint. The smoothing modifier serves to test the effect given by the shape and coarseness of a path. This modifier segmented paths and removed up to a specified number of segments in a section, while checking that this does not cause the path to intersect with land masses. The number of removals is upper bounded.

Intention-Based Anomaly Detection

We model ships as agents a_i and frequented way points, offshore platforms, and ports as goal sources s_j . We restricted the setting to a continuous 2D bounded region. Ships were modelled as changing speed negligibly in response to their attraction to goals, so the kinetic energy portion of the Lagrangian was ignored. Land masses and other ships in the scene were modelled as obstacle sources. While each goal was assigned its own ω_j , the ships and landmass points were restricted to share a single separate component for obstacles, $\omega_k = \omega_{obstacle} \forall k \in [R]$. This imposed that agents respond to obstacles in the same way whereas goals may vary in attractiveness and also accounted for the changing number of

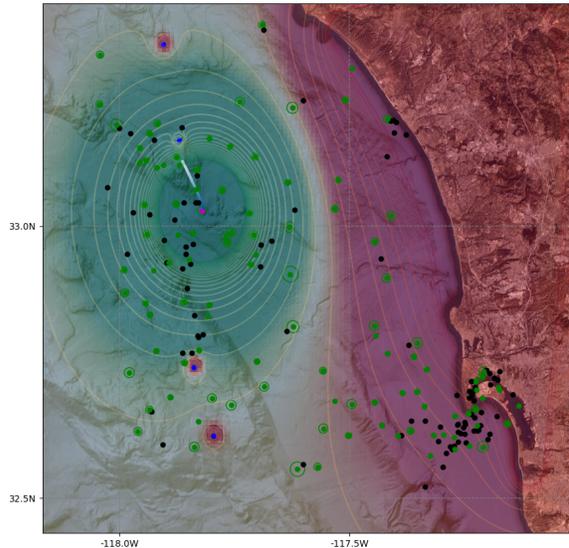


Figure 3: San Diego Heatmap and One Prediction

vessels to avoid.

To evaluate our algorithm, we compare it against other OC-SVMs. As a baseline, we compare against a Linear OC-SVM taking in the raw trajectories as input. We also compare against SVMs with different kernels, specifically, the polynomial and radial basis function (RBF) kernel. To compare these models, collections of real world AIS data belonging to four hour time windows were collected. Ship trajectories outside of a longitude and latitude bounding box were removed. A quarter of the windows were set aside for testing, and the rest were used for training. For each method, the mean of the training data latitudes and longitudes was also subtracted off.

To train our method, using the OC-SVM provided by sklearn, the set of structured features needed to be rescaled to avoid a positive feedback loop increasing the magnitude of the OC-SVM’s learned ω (Buitinck et al. 2013). This was used to account for SVM weight vector norm regularization. The feature was rescaled so that the largest L-1 norm of a feature in the dataset was 1.

Obstacle source locations were annotated for the scenes tested. Coastline and island polygons were obtained, and point sources were placed in an evenly spaced grid into regions within the polygons. Complete port annotations proved difficult to obtain. Instead, goal locations were sampled from trajectory data. 200 points representing potential goals were sampled randomly from the ship trajectories used. To train the baseline models, a grid search was performed to tune hyper parameters. These were the SVMs’ ν and error tolerance values, the polynomial svm’s degree, and the RBF-SVM’s γ value. For testing, synthetic AIS trajectories were added to the time windows set aside for testing. Models were evaluated based on their accuracy in detecting the artificially generated trajectories. In the computation of our features, the artificial trajectories were treated as ships in the scene and influenced repulsion.

		Long Beach				San Diego			
Mod.	Prop.	Intent	RBF	Poly	Linear	Intent	RBF	Poly	Linear
a)	2.5	61.2	54.3	49.5	52.7	49.9	60.8	58.9	53.2
	5	74.8	68.1	63.4	53.6	55.4	58.5	58.2	47.1
	10	86.1	80.4	62.5	48.7	61.6	63.5	56.7	47.0
	15	89.0	83.3	67.5	50.3	65.6	53.4	61.0	52.9
	20	88.6	84.6	66.9	50.7	79.6	51.6	54.2	52.0
b)	2.5	61.4	60.2	60.0	47.6	51.1	56.6	58.2	55.4
	5	65.0	65.3	60.1	55.1	45.5	55.6	55.8	51.2
	10	82.2	81.8	70.1	49.0	58.5	59.8	55.3	44.8
	15	83.8	81.9	73.8	48.5	69.0	66.5	61.6	51.0
	20	83.2	84.1	71.0	49.4	64.1	65.8	60.9	52.6
c)	2.5	60.5	59.9	58.4	46.9	57.1	55.4	56.5	63.5
	5	79.5	73.5	64.4	47.4	64.0	59.8	55.1	50.4
	10	93.0	85.0	65.2	49.8	65.9	51.1	53.4	58.7
	15	95.4	88.9	67.5	49.3	69.6	57.6	63.4	42.1
	20	95.2	84.5	71.5	49.6	80.9	60.6	51.7	54.5
d)	2.5	62.0	55.3	52.7	49.0	50.4	50.2	52.6	54.3
	5	73.4	67.3	62.8	50.3	56.5	52.8	59.2	45.8
	10	92.7	85.8	69.0	49.3	63.4	51.7	55.3	56.4
	15	94.3	89.3	74.4	49.7	72.7	52.4	52.3	51.2
	20	95.3	88.3	74.8	49.8	82.6	49.8	52.4	51.3
e)	2.5	61.5	57.0	50.5	56.8	54.7	59.7	58.4	50.4
	5	65.2	61.0	49.7	53.6	54.9	57.5	54.9	44.0
	10	77.6	74.2	54.5	54.1	60.3	58.7	58.2	52.3
	15	79.1	78.6	52.6	50.9	60.5	64.2	56.2	52.1
	20	80.3	75.5	54.5	51.6	64.0	65.8	55.2	55.3
30	82.5	80.6	54.1	52.4	68.0	72.3	59.3	49.3	

Table 1: The area under the curve scores for all comparative methods (in percentages). Scores are presented for both Long Beach and San Diego settings, for Faked trajectories with Modification types: a) All modifications included, b) No modification, c) smoothing, d) Divergence, and e) Loitering. Results are shown for test scenes with varied proportions of anomalous data (in percentages).

Experimental Results

Tables 1 and 2 summarize the performance of the different models in our anomaly prediction study. The intention-based detector outperforms the kernel SVM-based detectors in the Long Beach setting in accuracy of detecting our artificial trajectories. For only the unmodified artificial trajectories, our results are comparable but do not significantly outperform the RBF OC-SVM. This initial result shows promise for anomaly detection based on the intent-based formulation. Qualitatively, for the Long Beach scene, the coupled optimization appears to have the effect of selecting the highest priority goals from the many goals available. In figure 2 and 3 the goals corresponding to the attraction weights below 10% of the highest are marked in black, and the rest are marked in green. This shows a strong stratification in learned attraction weights based on popularity. Figure 1 plots known port locations from annotation. Comparing between the two images, many high weight goals match the known port locations. Some match the location of ports or offshore platforms not included in the known ports, such as at Catalina Island and the Beta Offshore Platform Ellen. There are also goals in open ocean reflecting common way points.

The San Diego setting included more open ocean and a greater variety of trajectories. All models had lower accuracy in this test scene. The simpler SVMs suffered from the lack of consistent routes in the setting which make it difficult

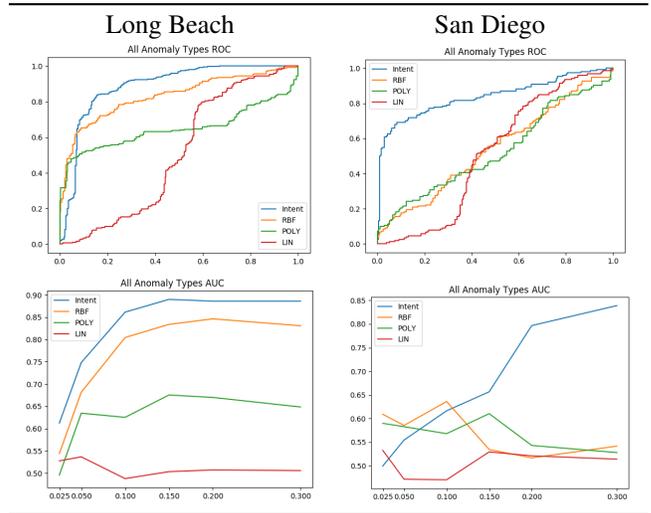


Table 2: AUC score by proportion of anomalous data in test scene and ROC curve for test scenes in which 30% of data are anomalous. Results are plotted for Intention prediction, Radial Basis Function Kernel, Polynomial Kernel and Linear (Raw) features in the Long Beach and San Diego settings.

to find support vectors in the raw data or in the kernels' resulting feature spaces. In our model, learned goal attractions are less stratified in this scene, reflecting the lack of consistent goals in the open ocean setting. Also, the goals in the ports and bays in the image do not appear to be prioritized either due to the large presence of ships following behaviors not captured by goal-based path planning. Our proposed model has higher accuracy scores for larger proportions of anomalies in test scenes, whereas the other OC-SVMs do not. Our conjecture is that the slight increase in robustness is due to the model characterizing some history of past trajectories as opposed to destinations in the spread out goal locations.

Conclusion

In this paper, we present a novel anomaly detection method for behaviors from trajectories. We evaluate the probability of a trajectory being anomalous based on prediction of agents' latent intentions, which, in this formulation are the sequences of goals and way points being pursued. Intention prediction is integrated with structured features based on path modelling using agent-based Lagrangian mechanics. Then, anomaly detection is carried out using a one-class support vector machine at the end of the learning framework. The prediction of intentions and training of the anomaly detector are coupled, allowing a single set of weights to be trained by alternating two optimizations. Experimental results distinguishing between faked trajectories in a marine surveillance show that our method is promising for detecting anomalous behaviors in the context of goal-directed agents. We hope that the method will contribute to the important applications of anomaly detection.

Acknowledgements

The work reported herein was supported by DARPA XAI grant N66001-17-2-4029, ONR MURI grant N00014-16-1-2007, ARO grant W911NF-18-1-0296, and SPAWAR grant N66001-17-2-3602.

References

- Agrawal, S., and Agrawal, J. 2015. Survey on anomaly detection using data mining techniques. *Procedia Computer Science* 60:708 – 713. Knowledge-Based and Intelligent Information & Engineering Systems 19th Annual Conference, KES-2015, Singapore, September 2015 Proceedings.
- Akoglu, L.; Tong, H.; and Koutra, D. 2015. Graph based anomaly detection and description: a survey. *Data Mining and Knowledge Discovery* 29(3):626–688.
- Alahi, A.; Goel, K.; Ramanathan, V.; Robicquet, A.; Fei-Fei, L.; and Savarese, S. 2016. Social lstm: Human trajectory prediction in crowded spaces. In *CVPR*.
- Buczak, A. L., and Guven, E. 2016. A survey of data mining and machine learning methods for cyber security intrusion detection. *IEEE Communications Surveys Tutorials* 18(2):1153–1176.
- Buitinck, L.; Louppe, G.; Blondel, M.; Pedregosa, F.; Mueller, A.; Grisel, O.; Niculae, V.; Prettenhofer, P.; Gramfort, A.; Grobler, J.; Layton, R.; VanderPlas, J.; Joly, A.; Holt, B.; and Varoquaux, G. 2013. API design for machine learning software: experiences from the scikit-learn project. In *ECML PKDD Workshop: Languages for Data Mining and Machine Learning*, 108–122.
- Bureau of Ocean Energy Management (BOEM) and National Oceanic and Atmospheric Administration (NOAA). Marinecadastre.gov. <http://www.marinecadastre.gov/data>. Retrieved: 2018-8-15.
- Chandola, V.; Banerjee, A.; and Kumar, V. 2009. Anomaly detection: A survey. *ACM computing surveys (CSUR)* 41(3):15.
- Chandola, V.; Banerjee, A.; and Kumar, V. 2012. Anomaly detection for discrete sequences: A survey. *IEEE Transactions on Knowledge and Data Engineering* 24(5):823–839.
- Görnitz, N.; Braun, M.; and Kloft, M. 2015. Hidden markov anomaly detection. In *ICML*, 1833–1842.
- Joachims, T.; Hofmann, T.; Yue, Y.; and Yu, C.-N. 2009. Predicting structured objects with support vector machines. *Communications of the ACM* 52(11):97–104.
- Kitani, K. M.; Ziebart, B. D.; Bagnell, J. A.; and Hebert, M. 2012. Activity forecasting. In *ECCV*.
- Kuderer, M.; Kretschmar, H.; Sprunk, C.; and Burgard, W. 2012. Feature-based prediction of trajectories for socially compliant navigation. In *RSS*.
- Li, K., and Fu, Y. 2014. Prediction of human activity by discovering temporal sequence patterns. *PAMI*.
- Li, W.; Mahadevan, V.; and Vasconcelos, N. 2014. Anomaly detection and localization in crowded scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36(1):18–32.
- Ma, W.-C.; Huang, D.-A.; Lee, N.; and Kitani, K. M. 2017. Forecasting interactive dynamics of pedestrians with fictitious play. In *CVPR*.
- Muller, K.-R.; Mika, S.; Ratsch, G.; Tsuda, K.; and Scholkopf, B. 2001. An introduction to kernel-based learning algorithms. *IEEE transactions on neural networks* 12(2):181–201.
- Nicholas Hoernle, Kobi Gal, B. G. P. P. A. R. 2018. Modeling the effects of students interactions with immersive simulations using markov switching systems. In *Educational Data Mining*, 292–298.
- Osekowska, E., and Carlsson, B. 2015. Learning maritime traffic rules using potential fields. In Corman, F.; Voß, S.; and Negenborn, R. R., eds., *Computational Logistics*, 298–312. Cham: Springer International Publishing.
- Qi, S.; Huang, S.; Wei, P.; and Zhu, S.-C. 2017. Predicting human activities using stochastic grammar. In *ICCV*.
- Qi, S.; Jia, B.; and Zhu, S.-C. 2018. Generalized earley parser: Bridging symbolic grammars and sequence data for future prediction. In *ICML*.
- Rhinehart, N., and Kitani, K. M. 2017. First-person activity forecasting with online inverse reinforcement learning. In *ICCV*.
- Rieck, K.; Krueger, T.; Brefeld, U.; and Mǎžller, K.-R. 2010. Approximate tree kernels. *Journal of Machine Learning Research* 11(Feb):555–580.
- Sidib, A., and Shu, G. 2017. Study of automatic anomalous behaviour detection techniques for maritime vessels. *Journal of Navigation* 70(4):847858.
- U.S. Department of Transportation (USDOT) / Bureau of Transportation Statistics (BTS). [usdot.opendata.arcgis.gov](http://www.usdot.gov). <http://geo.dot.gov/server/rest/services/NTAD/Ports/MapServer/0>. Retrieved: 2018-8-15.
- Vu, T.-H.; Olsson, C.; Laptev, I.; Oliva, A.; and Sivic, J. 2014. Predicting actions from static scenes. In *ECCV*.
- Walker, J.; Gupta, A.; and Hebert, M. 2014. Patch to the future: Unsupervised visual prediction. In *CVPR*.
- Wang, Z.; Deisenroth, M. P.; Amor, H. B.; Vogt, D.; Schölkopf, B.; and Peters, J. 2012. Probabilistic modeling of human movements for intention inference. *RSS*.
- Xie, D.; Shu, T.; Todorovic, S.; and Zhu, S.-C. 2018. Learning and inferring dark matter and predicting human intents and trajectories in videos. *PAMI* 40(7):1639–1652.
- Yuen, J., and Torralba, A. 2010. A data-driven approach for event prediction. In *ECCV*.
- Ziebart, B. D.; Ratliff, N.; Gallagher, G.; Mertz, C.; Peterson, K.; Bagnell, J. A.; Hebert, M.; Dey, A. K.; and Srinivasa, S. 2009. Planning-based prediction for pedestrians. In *IROS*.
- Zimek, A.; Schubert, E.; and Kriegel, H.-P. 2012. A survey on unsupervised outlier detection in high-dimensional numerical data. *Stat. Anal. Data Min.* 5(5):363–387.