

Adversary models account for imperfect crime data: Forecasting and planning against real-world poachers (Corrected Version)

Shahrzad Gholami¹, Sara Mc Carthy¹, Bistra Dilkina¹, Andrew Plumptre², Milind Tambe¹, Margaret Driciru³, Fred Wanyama³, Aggrey Rwetsiba³, Mustapha Nsubaga², Joshua Mabonga², Tom Okello³, Eric Enyel³

¹University of Southern California, {sgholami, saramarm, dilkina, tambe}@usc.edu,

²Wildlife Conservation Society, {aplumptre, mnsbuga, jmabonga}@wcs.org,

³Uganda Wildlife Authority, {margaret.driciru, fred.wanyama, aggrey.rwetsiba, tom.okello, eric.enyel}@ugandawildlife.org

ABSTRACT

Poachers are engaged in extinction level wholesale slaughter, so it is critical to harness historical data for predicting poachers' behavior. However, in these domains, data collected about adversarial actions are remarkably imperfect, where reported negative instances of crime may be mislabeled or uncertain. Unfortunately, past attempts to develop predictive and prescriptive models to address this problem suffer from shortcomings from a modeling perspective as well as in the implementability of their techniques. Most notably these models i) neglect the uncertainty in crime data, leading to inaccurate and biased predictions of adversary behavior, ii) use coarse-grained crime analysis and iii) do not provide a convincing evaluation as they only look at a single protected area. Additionally, they iv) proposed time-consuming techniques which cannot be directly integrated into low resource outposts. In this innovative application paper, we (I) introduce iWare-E a novel imperfect-observation aWare Ensemble (iWare-E) technique, which is designed to handle the uncertainty in crime information efficiently. This approach leads to superior accuracy and efficiency for adversary behavior prediction compared to the previous state-of-the-art. We also demonstrate the country-wide efficiency of the models and are the first to (II) evaluate our adversary behavioral model across different protected areas in Uganda, i.e., Murchison Fall and Queen Elizabeth National Park, (totaling about 7500 km²) as well as (III) on fine-grained temporal resolutions. Lastly, (IV) we provide a scalable planning algorithm to design fine-grained patrol routes for the rangers, which achieves up to 150% improvement in number of predicted attacks detected.

KEYWORDS

Predictive models; Wildlife poaching; Ensemble techniques; Field test evaluation; Wildlife protection

1 INTRODUCTION

Wildlife in Africa is currently under crisis, where animals such as elephants and rhinos are threatened by extreme poaching and habitat loss [20, 24]. Studies show that the elephant population has

decreased by 144,000 from 2007 to 2014, and continues to shrink by 8% each year continent-wide, primarily due to poaching activities for their ivory [3, 31]. Even though 84% of elephants currently reside in protected areas, they are still observed to have an extremely high rate of mortality [5] which serves to highlight the great need to take intelligent action towards thwarting poachers and reversing the downward trend in biodiversity loss.

Park rangers play a key role as the defenders of these protected areas, and are responsible for removing snares and traps placed by the poachers. Furthermore, they regularly collect records of illegal activities detected. While this data can provide significant insight and allow us to better model poachers' adversarial behavior, these records of attacks are unfortunately limited to the regions that the park rangers choose to visit (e.g., only about 60% of the protected areas are patrolled in each year). Moreover, the certainty about the absence of attacks largely depends on the amount of the patrol effort devoted to each area. Due to the vastness of the protected areas (e.g., Murchison Fall covers about 5000 sq. km shown in Figure 1), the limited number of outposts and rangers across the protected areas (e.g., about 30 outposts) and well-hidden placement of snares in the ground by poachers (Figure 2), it is not possible to conduct foot patrolling thoroughly throughout the area. Thus, it becomes necessary to consider this inherent uncertainty in real crime data in order to be able to use real data collected from the rangers we need to correctly model poachers' behavior.

Previous work on data-driven modeling of wildlife poachers' behavior suffers from the following limitations: (i) they learn poachers' behavior without reasoning about the corresponding uncertainty



Figure 1: Protected areas in Uganda: We present seasonal poachers' behavior analysis across two different protected areas (7500 sq. km in total). State-of-the-art focused only on a single area of 2500 sq. km with annual coarse-grained crime analysis.



Figure 2: Well-hidden snares detected by rangers, Photo credit: Uganda Wildlife Authority

in labels (due to insufficient amount of patrol effort)[16]. This results in unreliable predictions and consequently misleads the park rangers. Furthermore, (ii) they consider an annual basis for the temporal trend in crime predictions which results in missing short-term patterns in poachers' behavior [9, 16]. From a practical point of view, (iii) the computationally expensive techniques, including Markov Random Field and Dynamic Bayesian Networks [9, 26] proposed by many of these studies suffer from long runtimes and cannot be integrated into low resource outposts within the African protected areas. Last, to prove the reliability of the results to the law enforcement agencies in Uganda, models have to be evaluated in different sites. However, (iv) none of the previous studies showed their models' performance across multiple protected areas.

In this innovative application paper, we propose a new imperfect-observation **aWare Ensemble** (iWare-E¹) method which takes into account the major challenge of adversarial behavior modeling in the wildlife protection domain, i.e., imbalanced non-uniform uncertainty on evidence of crime collected by defenders. (I) This approach significantly improves accuracy and runtime of the algorithm (at least 90%) compared to state-of-the-art by using multiple fast running weak learners involved in a structured ensemble model compatible with the data collection scheme in protected areas. (II) we propose a scalable planning algorithm to design patrols, which utilizes the behavior prediction model (as a black box) and applies a piecewise linear approximation to reason about continuous values of patrol effort, which allows us to generate fine-grained patrols. We show that this approach results in up to 150% improvement in solution quality compared to the state-of-the-art. (III) Moreover, we evaluate all models on fine-grained temporal resolutions, i.e., seasonally, and for the first time, (IV) we evaluated all of our models on a larger scale based on real-world data across multiple protected areas including Murchison Fall and Queen Elizabeth in Uganda, covering 5000 sq. km and 2500 sq. km, respectively.

2 RELATED WORK

In data-driven wildlife protection literature, despite multiple effort to learn poachers' behavior from large scale, real historical data, previous work either does not consider the non-uniform uncertainty in data gathering by park rangers [16] or proposes time consuming techniques which cannot be directly integrated into computing systems available in low-resource outposts in the field [9, 26].

CAPTURE [26] was developed as a two-layered Bayesian Network with hidden variables to model imperfect detection of poaching. The main shortcoming with this approach was long runtime of the program which is a major obstacle to the deployment of the software. INTERCEPT [16] is a decision tree ensemble approach that assumes perfect detection of poaching activity by park rangers,

leading to biases in final predictions. Outperforming previous models, [9] proposed a hybrid model of Markov Random Field (MRF) and bagging ensemble of decision trees via a geo-clustering approach. This model selectively considers imperfect detection on some of the geo-cluster, however, still suffers from long runtime due to computationally expensive EM algorithms for parameter estimation. Unfortunately, other attempts to capture spatio-temporal patterns in illegal activity via Bayesian hierarchical models [6] did not report any standard metrics (e.g., precision and recall) to evaluate models' predictive performance.

Game theoretic models, in particular security games are well known to be effective models of protecting valuable targets against an adversary, and have been explored extensively at AAMAS [1, 14, 17, 18, 23] and the problem of patrol planning has been well studied in this context [2, 28]. However, much of this work assumes a perfectly rational adversary, which is not true for the wildlife protection domain, where poachers are boundedly rational. Green Security Games [8] were introduced to address the challenges specific to this domain, such as boundedly rational adversaries. While there has been work on learning these adversary models, this has been mostly done based on simulated games where data is collected by human subject experiments in the laboratory [10, 11, 14, 27, 34] rather than real world poachers. These methods are additionally unable to scale to real-world setups which typically have an enormous number of targets (e.g., 3900 targets of 1x1 sq. km in Murchison Fall park) and diverse geo-spatial characteristic.

In patrol planning for wildlife protection, PAWS was introduced as a risk-based randomized patrol generation algorithm which has been tested in the real world [7, 8, 33]. However, it relies on a specific type of explicit attacker behavior model such as Quantal Response and Subjective Utility Quantal Response [26]. Therefore, a framework for patrol planning to generate implementable patrolling routes against a black-box attacker was proposed in [32]. Although this framework can handle complex data-driven predictive model, it was not able to scale up for continuous patrol effort values. Furthermore, none of the aforementioned studies account for naturally occurring uncertainty in crime evidence collected by defenders and its consequent effects on planning.

In ensemble modeling literature, ensemble-based techniques are well-known to improve performance of single models (i.e., weak learner) and they have been widely used to address imbalance in positive and negative instances of observations in a variety of domains from chaotic behavior modeling for stock market prediction [4], knowledge base population in text analysis [29] and vowel discrimination tasks [13]. Ensemble techniques can be categorized as iterative based ensembles or parallel ensembles [12]. In adversary behavior modeling [16] leverages iterative based ensembles. However, parallel ensembles which are based on parallel re-sampling and bagging of weak learners have also been shown to be very time saving and easy to develop in many practical problems to learn human behavior, e.g., in online banking fraud detection [30]. In this paper, we propose a parallel ensemble for which we re-sample via filtering of negative instances of crime depending on the amount of the defenders' effort to collect those instances. By this we are able to minimize the adverse effects of uncertainty in negative instances of crime and boost the prediction accuracy by generating more specialized weak learners based on more confident subsets of data.

¹To be pronounced similar to ivory

3 PREDICTIVE MODEL AND ALGORITHM

3.1 Domain Features

The wildlife crime datasets in this paper are from Uganda. We study Murchison Fall National Park jointly with Bugungu and Karuma wildlife reserves, and Queen Elizabeth National Park with Kigezi and Kyambura wildlife reserves. We refer to these protected areas as MFPA and QEPA, which span about 5000 sq. km and 2500 sq. km, respectively. There are 30 and 20 patrol posts situated across these protected areas from which Uganda Wildlife Authority (UWA) rangers conduct patrols. Along with the amount of patrolling effort in each area, the datasets contain 14 years (2003-2016) of the type, location, and date of wildlife crime activities. To study wildlife crime, we divide the protected areas into 1 sq. km grid cells. Each of these cells is associated with several static geo-spatial features such as terrain (e.g., slope), distance values (e.g., distance to border, roads, and towns), and animal density. Additionally, each cell is associated with dynamic features such as patrol effort (coverage) across time and observed illegal activities (e.g., snares). Patrol effort is the amount of distance walked by park rangers across a cell at a specific time step. Since park rangers do not have unlimited manpower to patrol each cell thoroughly, it is possible that the amount of distance walked by them is not sufficient and consequently, some of the well-hidden snares are not detected by them. This fact is the source of uncertainty over the negative instances of crime and has to be considered in the adversarial reasoning.

3.2 Dataset Preparation

We create the wildlife crime datasets, $\mathcal{D} = (X, y, w)$, studied in this paper from a dataset of recorded illegal activity by discretizing the records by time and by location so that we have a set of T time steps and N locations. $X \in \mathbb{R}^{TN \times f}$ is a matrix of f predictor features recorded at each of these T discrete time steps and N locations. Each row of predictor features $X(k)$ includes several time-invariant geo-spatial features (discussed earlier) associated with each location (e.g., average animal density, slope, forest cover, net primary productivity, distance from patrol post, town, rivers, park boundaries, salt licks and water holes) and a set of time-variant covariates, patrol effort $c_{t-1}(k)$, that is the amount of patrol coverage during the previous time step $t - 1$, which models the potential deterrence effect of patrols and $c_t(k)$ the amount of patrol effort in the current time step, which models the effort devoted to each data point at the data collection time. $y \in \{0, 1\}^{TN}$ denotes the observation vector associated with all data points. Additionally, each data point in the dataset is associated with a weight $w \in \{0, 1\}^{TN}$. In the original dataset all weights are 1, however, if data point k is recognized as a sufficiently uncertain data point by the algorithm, $w(k)$ will be changed to 0 and k is disregarded from the training set. To train any predictive model in this study, we divide this data into two sets for training, \mathcal{D}^{tr} , and testing, \mathcal{D}^{ts} . For our study, we used a training set which includes the first $T - 1$ years of crime data (corresponding to 6 years) and tested on the data in next successive year.

3.3 Uncertainty in Poaching Activity Detection

While park rangers attempt to remove and record any illegal activity signs (e.g., snares and traps), it is often the case that they do not

detect such signs, particularly if the snares are well-hidden. The success with which they detect these signs is linked to the amount of effort exerted in patrolling these regions. While positive records of poaching are assumed to be reliable in this study regardless of the amount of patrol effort, there is an intrinsic uncertainty associated with negative labels in the dataset, which depends on the patrol effort amount c_t (i.e., distance walked) devoted to each region during the data collection period, t . In particular given a threshold for patrolling effort θ , negative data samples recorded based on a patrol effort of $c_t \geq \theta$ are relatively more reliable (i.e., more probable to be actual negative samples) compared to the ones that were recorded based on a patrol effort of $c_t \leq \theta$. We use the notation subscript of θ_i^- to represent an instantiation of weight vectors in our dataset where negative samples recorded by a patrol effort of $c_t \leq \theta_i$ are ignored. In other words, for each data point k in \mathcal{D} , if $y(k) = 1$, then $w_{\theta_i^-}(k) = 1$. If $y(k) = 0$ then $w_{\theta_i^-}(k) = 1$ when $c_t(k) \geq \theta_i$ and $w_{\theta_i^-}(k) = 0$ when $c_t(k) \leq \theta_i$.

3.4 Imperfect Observation-aware Ensemble

Due to diversified and robust characteristics of ensemble techniques, we propose a new imperfect observation-aware Ensemble model (iWare-E), which is able to handle the intrinsic uncertainty in the poaching activity data collection scheme by park rangers mentioned earlier. This ensemble technique outlined in Algorithm 1 involves multiple weak learners (also known as experts or ensemble members) which vote on the final predictions. Each weak learner is trained based on a subset of the dataset, $\mathcal{D}_{\theta_i^-}$, filtered by a threshold θ_i where i is in $\{0, 1, \dots, I - 1\}$ and $\theta_i \leq \theta_{i+1}$. Line 2 in Algorithm 1 indicates that for any choice of $[\theta_{min}, \theta_{max}]$, I number of equally or unequally distanced intermediate thresholds θ_i can be obtained such that $\theta_{min} \leq \theta_i \leq \theta_{max}$ and consequently I weak learners, $C_{\theta_i^-}$ can be trained on the corresponding $\mathcal{D}_{\theta_i^-}$ (line 6 in Algorithm 1). Figure 3(a) shows how patrol effort is filtered by different thresholds to generate a different sub-dataset and a corresponding expert in ensemble. The leftmost branch in the Figure 3(a) represents the case that $\theta_0 = 0$, i.e., the entire dataset and the rightmost branch represents the case where negative instances of crime associated with $c_t \leq 2$ are disregarded.

To address the voting scheme among the ensemble members, $C_{\theta_i^-}$, we propose a binary vote qualification matrix, V^q which determines the qualification, 1, or disqualification, 0, of weak learners (each represented by a row), across ranges of c_t indicated by $[\theta_i, \theta_{i+1})$ (each represented by a column). Since each of these models are qualified to make predictions on data points which fulfill the condition $c_t \geq \theta_i$, the vote qualification matrix is a triangular matrix with size $I \times I$ (lines 7 through 14 in Algorithm 1). An example of this matrix is illustrated via the table in Figure 3(b), where each column represents an interval on $c(t)$ and each row represents a trained expert in the ensemble. It is worth noting that number intervals and number of experts are always equal (denoted by I here). If an expert is qualified to make predictions on an interval, the corresponding V^q element is 1. Furthermore, we also introduce a vote power matrix V^p of size $I \times I$ which contains the weights or vote power of each of the weak learners (each represented by a row), across ranges of c_t indicated by $[\theta_i, \theta_{i+1})$ (each represented by a column). An example of a vote power matrix is shown with

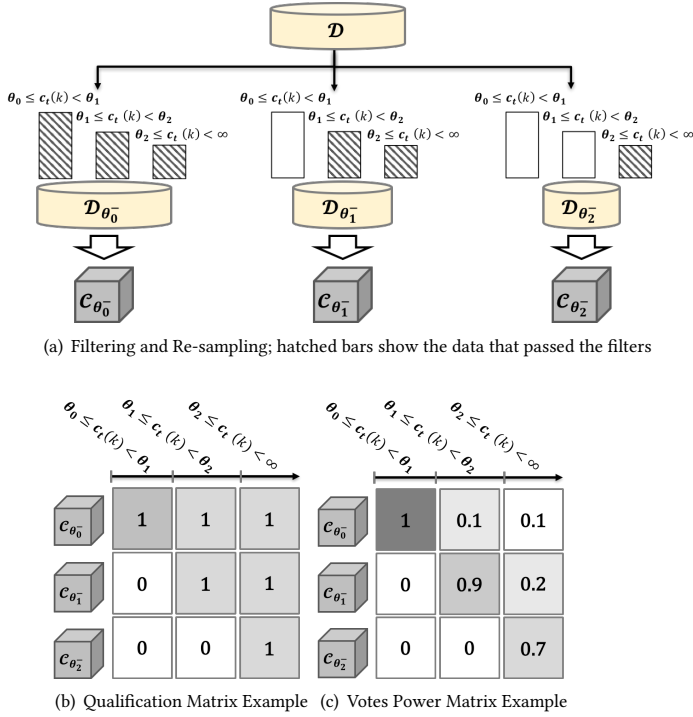


Figure 3: Schema of iWare-E model

different shade of gray rectangles and numbers associated with them in Figure 3(c).

The actual weights on the weak learners are a combination of qualification and vote power matrices, $\mathbf{V}^{qp} = \mathbf{V}^q \circ \mathbf{V}^p$. To ensure proper weighing of qualified weak learners within each range of $[\theta_i, \theta_{i+1})$, \mathbf{V}^{qp} is normalized such that each column sums up to one (lines 16 and 17 in Algorithm 1). While \mathbf{V}^q depends on the structure of the ensemble method, \mathbf{V}^p is a hyper parameter. To tune this hyper parameter, we choose an initial \mathbf{V}_0^p and a validation set, and then we use Algorithm 1 to minimize the error between actual observations and estimations by the model. This tuned \mathbf{V}^p is used for training the ensemble on other sets via Algorithm 1. To make prediction on the test set and evaluate the model, the appropriate interval (which depends on the value of c_t) is obtained (line 3 in Algorithm 2) and then, the weighted average of all experts' predictions is computed using $\bar{\mathbf{V}}^{qp}$ (line 3 and 7 in Algorithm 2).

4 MODEL EVALUATION

4.1 Evaluation on Historical Data

For the illegal activity datasets we use, each protected area is divided into small 1×1 km regions and time steps of 3 months long are considered as opposed to the state-of-the-art [9] that considered coarse time steps of one year long, which makes it vulnerable to missing fine-grained temporal trends in poaching. To convince law enforcement agencies, it was essential to evaluate the predictive model across different protected areas and demonstrate superior performance of the model for smaller temporal resolutions.

Algorithm 1: Train iWare-E

input : Train dataset (\mathcal{D}^{tr} , \mathbf{w}^{tr} , \mathbf{c}_t^{tr});
 Threshold parameters (θ_{min} , θ_{max} , I);
 Vote power matrix, (\mathbf{V}^p , size $I \times I$)

output: Classifiers and weights matrix ($C_{\theta_i^-}$ and $\bar{\mathbf{V}}^{qp}$)

- 1 find threshold values for I intervals on c_t ;
- 2 $\theta \leftarrow \text{FindThresholdVector}(\theta_{min}, \theta_{max}, I)$;
- 3 train the classifiers;
- 4 **for** $i \leftarrow 0$ **to** $I - 1$ **do**
- 5 $\mathcal{D}^{tr}, \mathbf{w}_{\theta_i^-}^{tr} \leftarrow \text{FilterData}(\mathcal{D}^{tr}, \mathbf{w}^{tr}, \mathbf{c}_t^{tr})$;
- 6 $C_{\theta_i^-} \leftarrow \text{TrainABaggingEnsemble}(\mathcal{D}^{tr}, \mathbf{w}_{\theta_i^-}^{tr})$;
- 7 build vote qualification matrix, row is a member and column is an interval on c_t ;
- 8 **for** $j \leftarrow i$ **to** $I - 1$ **do**
- 9 $\mathbf{V}^q(j, i) \leftarrow 1$;
- 10 **end**
- 11 **for** $k \leftarrow 0$ **to** $i - 1$ **do**
- 12 $\mathbf{V}^q(k, i) \leftarrow 0$;
- 13 **end**
- 14 **end**
- 15 find total weights for member;
- 16 $\mathbf{V}^{qp} \leftarrow \text{MultiplyElementWise}(\mathbf{V}^q, \mathbf{V}^p)$;
- 17 $\bar{\mathbf{V}}^{qp} \leftarrow \text{ColumnWiseNormalizeToSumOne}(\mathbf{V}^{qp})$;

Algorithm 2: Predict by iWareE

input : Test dataset (\mathcal{D}^{ts} , \mathbf{w}^{ts} , \mathbf{c}_t^{ts});
 Threshold parameters (θ_{min} , θ_{max} , I);
 Classifiers and weights matrix ($C_{\theta_i^-}$ and $\bar{\mathbf{V}}^{qp}$)

output: Predicted probability of crime observation (\mathbf{p})

- 1 test the classifiers;
- 2 **for** $\mathcal{D}^{ts}(k) \in \mathcal{D}^{ts}$ **do**
- 3 $i^* \leftarrow \text{FindRelatedInterval}(\mathbf{c}_t^{ts}(k))$;
- 4 **for** $i \leftarrow 0$ **to** $I - 1$ **do**
- 5 $\mathbf{p}(k) \leftarrow \mathbf{p}(k) + C_{\theta_i^-}(k) \cdot \bar{\mathbf{V}}^{qp}(i, i^*)$;
- 6 **end**
- 7 **end**

For these datasets, the patrol effort is the amount of distance that park rangers walk through a 1×1 km region during a single time step of study. We tune hyper parameter based on training from 2007-2012 and validating on the 2013 dataset. Three different sets are used to evaluate our model, trained on the data from the years 2008-2013, 2009-2014 and 2010-2015 and tested on 2014, 2015 and 2016 respectively. Due to space consideration, detailed comparison of the proposed model with all possible baselines (e.g., Positive, Random, Training Label baselines) are presented in the supplementary material in the online Appendix². We selected $\theta_0 = 0$ and $\theta_{I-1} = 7.5$ with 16 equally-distanced intermediate values of θ_i . Since the number of the data points with $c_t > 7.5$ was significantly

²<https://www.dropbox.com/s/cu08xr0txd8ur41/Appendix.pdf?dl=0>

Table 1: Comparing all models' performances for MFPA

Test	2016					
	state-of-the-art				iWare-E	
Mdl.	SVB	DTB	MRF	HY	SVB-iW	DTB-iW
AUC	0.59	0.60	0.63	0.69	0.64	0.68
Prec.	0.20	0.17	0.19	0.22	0.20	0.21
Recall	0.40	0.57	0.59	0.62	0.57	0.58
F1	0.27	0.26	0.28	0.32	0.29	0.30
L&L	0.61	0.73	0.82	1.01	0.89	0.89
L&L %	8.73	10.41	11.02	14.4	11.0	12.0
Test	2015					
	state-of-the-art				iWare-E	
Mdl.	SVB	DTB	MRF	HY	SVB-iW	DTB-iW
AUC	0.58	0.63	0.67	0.69	0.68	0.69
Prec.	0.18	0.17	0.19	0.21	0.20	0.19
Recall	0.41	0.59	0.62	0.61	0.62	0.65
F1	0.25	0.26	0.29	0.31	0.30	0.30
L&L	0.59	0.80	0.95	1.02	0.99	1.01
L&L %	8.43	11.37	11.95	14.52	12.0	13.0
Test	2014					
	state-of-the-art				iWare-E	
Mdl.	SVB	DTB	MRF	HY	SVB-iW	DTB-iW
AUC	0.56	0.57	0.68	0.66	0.68	0.69
Prec.	0.23	0.19	0.26	0.26	0.25	0.26
Recall	0.33	0.54	0.64	0.61	0.62	0.64
F1	0.27	0.28	0.37	0.36	0.36	0.37
L&L	0.46	0.62	1	0.96	0.96	1.02
L&L %	7.68	10.36	16.3	15.97	16	17

lower compared to the ones with $c_t \leq 7.5$, we chose $\theta_{max} = 7.5$ to guarantee reasonable training datasets for all weak learners.

We compare the performance of the proposed model with the latest best performing existing models examined on the QEPA dataset in [9] in terms of standard machine learning metrics including AUC, Precision, Recall, F1. Since the metrics are used to evaluate models on datasets with no uncertainty in the underlying ground truth, we also use the L&L metric [19], which is a metric specifically designed for models learned on Positive and Unlabeled datasets. L&L is defined as $L\&L = \frac{r^2}{Pr[f(Te)=1]}$, where r denotes the recall and $Pr[f(Te) = 1]$ denotes the probability of a classifier f making a positive class label prediction and is estimated by the percentage of positive predictions made by the model on a given test set. We also discuss our algorithm runtime compared to the state-of-the-arts.

Table 1 summarizes the performance of different models including bagging ensemble of SVM (denoted as SVB), bagging ensemble of decision trees (denoted as DTB), Markov Random Field (denoted as MRF), hybrid of last two ones (HY) (presented in [9]) as the existing models in literature against iWare-E model with two different weak learners including SVB and DTB, which are denoted as SVB-iW and DTB-iW, respectively. For all cases, applying iWare-E ensemble on SVB or DTB improves the AUC and L&L scores up to 10% compared to SVB or DTB. Furthermore, in comparison with the HY model, which is the state-of-the-art model, our results are improved up to 10% (see recall for QEPA 2014 test set) while the runtime is significantly reduced as we will discuss later.

Table 2: Comparing all models' performances for QEPA

Test	2016					
	state-of-the-art				iWare-E	
Mdl.	SVB	DTB	MRF	HY	SVB-iW	DTB-iW
AUC	0.53	0.61	0.58	0.68	0.69	0.71
Prec.	0.13	0.08	0.08	0.11	0.13	0.12
Recall	0.13	0.59	0.58	0.62	0.57	0.62
F1	0.13	0.14	0.14	0.18	0.21	0.19
L&L	0.31	0.82	0.79	1.19	1.28	1.27
L&L %	1.85	4.84	4.43	7	7.0	7.0
Test	2015					
	state-of-the-art				iWare-E	
Mdl.	SVB	DTB	MRF	HY	SVB-iW	DTB-iW
AUC	0.53	0.63	0.62	0.7	0.66	0.71
Prec.	0.12	0.09	0.09	0.1	0.12	0.12
Recall	0.09	0.60	0.59	0.62	0.47	0.61
F1	0.10	0.16	0.15	0.18	0.19	0.2
L&L	0.18	0.90	0.81	1.06	0.94	1.21
L&L %	1.12	5.62	4.98	6.6	6.0	7.0
Test	2014					
	state-of-the-art				iWare-E	
Mdl.	SVB	DTB	MRF	HY	SVB-iW	DTB-iW
AUC	0.58	0.70	0.68	0.74	0.68	0.76
Prec.	0.12	0.07	0.06	0.08	0.09	0.09
Recall	0.17	0.65	0.62	0.62	0.52	0.72
F1	0.14	0.12	0.11	0.14	0.15	0.16
L&L	0.53	1.17	1.02	1.34	1.22	1.75
L&L %	1.95	4.33	3.76	4.95	5.0	6.0

Table 3: Average runtime over all years (seconds)

Mdl.	state-of-the-art				iWare-E	
	SVB	DTB	MRF	HY	SVB-iW	DTB-iW
MFPA	80	65	45850	15328	1250	183
QEPA	78	71	31115	10348	1309	175

Currently, the SMART software (a Spatial Monitoring and Reporting Tool) is used worldwide by park rangers to collect data and make decisions about patrolling routes. However, this software does not exploit historical data to predict poachers' behavior. In order to make possible the integration of our software in platforms like SMART, we have to develop fast models able to be run on non-high-performance machines. To that end, we present runtime analysis of all models for both parks in Table 3. Notably, DTB-iW completes in less than 200 seconds and SWD-iW completes in about 1300 seconds, which are significantly lower compared to MRF and HY models that suffer from the slow speed of EM algorithm. Although, SVB and DTB are fastest, they do not perform well in terms of accuracy as discussed earlier. All the experiments were performed on a machine with 2.6GHz and 8GB RAM.

4.2 Field Tests Results

Fortunately, for QEPA, we have access to the field test data studied in [9], where attack prediction labels were defined as the proportion of 1×1 sq. km cells inside a patrol area of 3×3 sq. km, that were

Table 4: Comparison across sources of prediction, QEPA

Months		11/1-1/31			2/1-4/30			5/1-6/30		
Counts		Ar	C	E	Ar	C	E	Ar	C	E
HY	H	5	1	10.1	5	9	8.9	5	7	10.51
	L	22	0	11	22	3	8.87	22	2	10.3
iW	H	4	1	9.95	9	10	9.1	10	9	9.45
	L	23	0	10	18	2	9.83	17	0	10.36

Table 5: Statistical Significance Test: p values (<0.05)

Months	11/1-1/31	2/1-4/30	5/1-6/30
HY	2.74×10^{-4}	5.76×10^{-4}	3.22×10^{-3}
iW	2.24×10^{-3}	8.83×10^{-5}	1.24×10^{-3}

predicted to be attacked by the model. When attack prediction rate was more than 50%, the area was classified as high (H) and when it was less than 50%, the area was classified as low (L). So two experiment groups of high and low were generated according to model's attack prediction rates from November 2016 - June 2017. Table 4 summarizes the field test results for the QEPA dataset across time and prediction sources. The first row indicates three time steps of three months long when field tests were executed by park rangers. For each of these time steps and for each source of prediction, three different values are reported for high (H) and low (L) regions, i.e., Ar, C, E denote the number of 3×3 sq. km regions, counts of observations and amount of patrol effort, respectively. Predictive model names are outlined in the first column. Hybrid model and iWare-E models are indicated by HY and iW, respectively. iW denotes the best performing model for QEPA which is DTB-iW.

Predictions of the iWare-E model depend on the amount of patrol effort devoted to each cell which is designed to take into account the detectability challenge. We assumed a fixed amount of 4.5 km for patrol effort during 3 months for each individual cell. This value was selected based on the historical distribution of patrol effort over different months. Table 4 demonstrates that iWare-E model is more selective between high and low groups and outperforms state-of-the-art predictions (HY). Furthermore, average monthly values of patrol effort, column named as E, computed for each high and low groups shows that park rangers are covering areas nearly uniformly. This indicates the shortcoming of their current method for planning as poachers are not attracted to all regions equally. Table 5 summarizes the statistical significance test based on occurrence of attack for all sources of prediction across time. Chi-Square test results for binary outcomes show that for both predictive models at all time steps, p-value is below the significance level (0.05), which determines that it is not likely that observations are due to chance.

5 PATROL PLANNING

The goal of developing these predictive models is to allow the rangers to leverage this additional information in order to better detect and reduce the number of attacks in protected areas. While there has been much work in Green Security Games (GSG) doing patrol planning in these domains [8, 21], much of this work has assumed explicit models for how poachers behave. These models, ranging from perfect rationality to bounded rationality models like

Quantal Response (QR) and Subjective Utility Quantal Response (SUQR) [22, 25] can make planning much simpler due to their explicit nature. However when there is access to data on poaching activity, we can achieve much more accurate representations poachers behavior with machine learning models. These models are much more difficult to optimize from a planning perspective since we only have black box access to the predictions given a desired input. While there has been some prior work in GSG planning patrols which optimize black box functions [32], which we build off of, there are several key differences which make it so that these solution methods are not appropriate for our problem. The most important is that previous work is limited to optimizing over discrete levels of patrol effort. For more general machine learning models such as iWare-E, which can make predictions based off of continuous values of patrol effort, this can result in either large losses in solution quality when discretization levels are too coarse, or large runtimes when discretization is too fine (which we show in Figures 5 and 6). To address these issues we propose to instead approximate the machine learning model through the use of piecewise linear (PWL) functions. This allows us to reason about continuous values of patrol effort and achieve significant improvements in solution quality (up to 150% improvement) while remaining scalable (up to 400x increase in speed).

Following standard practice in Green Security Games [8, 15] we model the wildlife conservation patrolling problem as a game played on a graph $G(N, E)$ of nodes and edges, over a period of time T . We discretize the conservation area into a set of N grid cells, corresponding to the 1×1 km regions of the dataset. In order to protect the conservation area, rangers conduct patrols over these N grid cells. Patrolling a grid cell takes a certain amount of time and effort, and we assume that the ranger may only spend T time steps patrolling in any given day. Note that this time discretization is distinct from the 3 month long time steps considered in the dataset. Here we define a time step as the minimum amount of time it would take to cross a single grid cell (so that the ranger must spend at least 1 time step in each grid cell they choose to visit). A single patrol corresponds to a 1 unit flow on the time unrolled graph $G(N', E')$, with a set of nodes and directed edges given by:

$$N' := \{v' = (v, t) : v \in N \ t \in \{1, T\}\}.$$

$$E' := \left\{ \begin{array}{l} (u, v) \in E \cup \{(w, w) \mid u, v, w \in N\} \\ ((u, t_1), (v, t_2)) : \begin{array}{l} t_2 = t_1 + 1 \\ t_1, t_2 \in \{1, T\} \end{array} \end{array} \right\}.$$

One of the grid cells is a designated patrol post. All patrols must begin and end at this grid cell, and so we designate this grid cell as the source $s \in N$. For notational convenience, let $s_1 = (s, 1) \in N'$ and $s_T = (s, T) \in N'$ the source node in the time unrolled graph at the first and last time steps respectively. The goal of these patrols is to detect signs of poaching activity, and rangers may conduct multiple rounds of successive patrols. The ability of a patrol to detect illegal activity at any grid cell $v \in N$ will depend on the level of patrol effort at that cell c_v^r (where r indicates the r^{th} round of patrols). Each patrol may expend a total of T units of patrol effort on any single patrol. Since a single unit of patrol effort is necessary to cover any cell for a single time step, a feasible patrol corresponds to a single unit flow on G' originating at the source s_1 and where the sum of the total flow on all edges in E' is equal to T . We then




Figure 4: Sample heat map of patrol efforts computed for the QEPA (left) and MFPA(right)

denote the set of feasible patrol efforts as \mathcal{F} given by:

$$\mathcal{F} := \left\{ f_{u',v'} : \begin{array}{l} \sum_{u:(u',v') \in E'} f_{u',v'} = \sum_{u':(v',u') \in E'} f_{v',u'} \quad \forall v' \in N' \\ \sum_{u':(s_1,u') \in E'} f_{s_1,u'} = \sum_{u':(u',s_T) \in E'} f_{u',s_T} = 1 \\ \sum_{(v',u') \in E'} f_{u',v'} = T \end{array} \right\}$$

Optimizing Detection Probability: We assume that for each grid cell there exists function $g_v : C_v^r \times C_v^{r-1} \rightarrow P_v^r$ which maps the current and past total defender patrol effort at a particular grid cell $v \in N$ to a corresponding likelihood that there will be a detected attack P_v^r at that grid cell in round r . What we would like to do is solve for a series of patrols which maximizes the probability of detecting attacks over the entire area. The rangers conduct a total of K patrols within a single round r ; since each patrol has a probability $\sum_{u':(u',v') \in E'} f_{u',v'}$ of visiting each cell v (ie. the sum of the total flow visiting that cell across all time steps), the expected aggregate patrol effort c_v at v is this probability times the total number of patrols K . The following Mathematical Program (MP) computes the optimal patrol effort which maximizes the predicted total detected attacks:

$$\begin{aligned} \max_{c,f} \quad & \sum_{v \in N} g_v(c_v^r, c_v^{r-1}) \\ & f_{u',v'} \in \mathcal{F} \quad \forall (u',v') \in E' \\ & K \sum_{u':(v',u') \in E'} f_{u',v'} = c_v \quad \forall v \in N, v' = (v,t) \\ & \sum_{v \in N} c_v = T \times K \end{aligned} \quad (\mathcal{P})$$

Using the iWare-E model to generate these predictions we only have black box access to these functions g_v , so we instead use piecewise linear (PWL) approximations of the g_v in our objective functions. In order to construct these functions, we build datasets D_g of $m_r \times m_{r-1} \times N$ sample points p from the N functions g_v , giving the probability of detection P_v for m_r possible effort values for the current round and m_{r-1} effort values for the previous rounds of patrol:

$$D_g := \left\{ p = \langle C_{v,i}^r, C_{v,j}^{r-1}, g_v(C_{v,i}^r, C_{v,j}^{r-1}) \rangle : \begin{array}{l} \forall v \in N \\ i \in \{1, m_r\} \\ j \in \{1, m_{r-1}\} \end{array} \right\}$$

Using this dataset D_g we can construct our PWL approximation by representing any set of feasible patrol efforts (c^r, c^{r-1}) and corresponding predicted detection of attack $g(c^r, c^{r-1})$ as a convex combination of their nearest neighbors in the dataset D_g .

At round r we already have data on the previous $r-1$ round's patrolling effort at each cell which we denote \tilde{c}_v^{r-1} . We use the notation capital $C \subset p \in D_g$ to denote patrol effort data used to construct the piecewise linear objective and lowercase \tilde{c} to denote known past patrol effort data. Because the \tilde{c}_v^{r-1} are known we can directly express them as a convex combinations of the closest two data points (C^+, C^-) so that $\tilde{c}_v^{r-1} = \lambda_v^{r-1} C_v^+ + (1 - \lambda_v^{r-1}) C_v^- \quad \forall v \in N$. We want to plan patrols for the current round r , meaning that for the patrol effort c_v^r , we do not know beforehand what the two

Segments	Runtime (s)		Detections		Error (%)		
	QEPA	MFPA	QEPA	MFPA	QEPA	MFPA	
(\mathcal{P}_1)	5	0.25	0.26	8.8	18.8	17	2.7
	10	1.9	1.8	14.5	22.7	9.4	2.9
	20	1.2	12.2	17.4	23.5	0.2	8.2
	40	20.2	57.1	19.4	25.5	4.3	2.7
	80	45.5	97.9	19.7	26.7	4.2	1.8
(\mathcal{P}_2)	25	21.4	4.3	11.1	22.5	37.7	11.1
	36	131.7	104	13.4	23.8	15.4	7.7

Table 6: Performance of the PWL approx. MILP ((\mathcal{P}_1)) (top) and 2D-PWL approximation MILP ((\mathcal{P}_2)) (bottom) with increasing segments.

closest data points in D_g will be. Instead we express c_v^r as a convex combination of *all* points p_c , and constrain the weights λ^r on the points to belong to a *Specially Ordered Set of Type 2 (SOS2)* which are an ordered set of variables where at most two consecutive variables may be non-negative. The objective function of MP (\mathcal{P}) can then be expressed as:

$$\begin{aligned} g_v(c_v^r, \tilde{c}_v^{r-1}) &= \sum_i \lambda_{v,i}^r \left(\lambda_v^{r-1} g_v(C_{v,i}^r, C_v^+) \right. \\ &\quad \left. + (1 - \lambda_v^{r-1}) g_v(C_{v,i}^r, C_v^-) \right) \\ &= \sum_i \lambda_{v,i}^r \tilde{g}_v(C_v^r) \end{aligned} \quad (1)$$

Where we add additional constraints:

$$\begin{aligned} \sum_{i \in [m_r]} \lambda_{v,i}^r C_{v,i}^r &= c_v^r \quad \forall v \in N \\ \lambda_{v,i}^r &\in \text{SOS2} \quad \forall v \in N, i \in [m_r] \\ \lambda_{v,i}^r &\in [0, 1] \quad \forall v \in N, i \in [m_r] \\ \sum_{i \in [m_r]} \lambda_{v,i}^r &= 1 \quad \forall v \in N \end{aligned} \quad (2)$$

So that MP (\mathcal{P}) is now expressible as a Mixed Integer Linear Program which we refer to as MILP (\mathcal{P}_1).

Two Stage Planning: Given that the predictions are functions of past and current patrol, we have the ability to plan for multiple rounds of patrolling. We want to maximize the probability of detecting an attack in two rounds, r and $r+1$, ie. $\sum_{v \in N} g_v(c_{v,i}^r, \tilde{c}_v^{r-1}) + \sum_{v \in N} g_v(c_{v,i}^{r+1}, c_v^r)$. We already know how to construct the PWL approximation of $g_v(c_{v,i}^r, c_v^{r-1})$ since we have \tilde{c}_v^{r-1} as data; however both c_v^r and c_v^{r+1} are variables and must be expressed as convex combinations of points in \mathcal{D} using the same type constraints as (2). We can then express the tuple (c_v^r, c_v^{r+1}) as a convex combination of the 4 closest points in \mathcal{D} with weights $\Lambda_{i,j}$ using from the weights $\lambda_{v,i}^r$ and $\lambda_{v,j}^{r+1}$ with the following constraints $\sum_i \Lambda_{i,j}^v = \lambda_{v,j}^{r+1} \quad \forall v \in N, j \in [m_{r+1}]$ and $\sum_j \Lambda_{i,j}^v = \lambda_{v,i}^r \quad \forall v \in N, i \in [m_r]$. With these we are guaranteed to have only 4 non-zero $\Lambda_{i,j}^v$ since there are only 2 non-zero $\lambda_{v,i}^r$ and $\lambda_{v,j}^{r+1}$. The two stage optimization problem, MILP

(\mathcal{P}_2) can then be expressed as:

$$\begin{aligned}
\max_{\lambda, \Lambda} \quad & \sum_{v \in N} \sum_{i,j} \Lambda_{i,j}^v g(C_{v,i}^{r+1}, C_{v,j}^r) + \sum_i \lambda_{v,i}^r \tilde{g}_v(C_{v,i}^r) \\
f_{u',v'}^r, f_{u',v'}^{r+1} \in \mathcal{F} \quad & \forall (u', v') \in E' \\
K \sum_{u':(v',u') \in E'} f_{u',v'}^r = c_{v'}^r \quad & \forall v \in N, v' = (v, t) \\
\sum_{v \in N} c_v^r = T \times K \\
\sum_i \lambda_{v,i}^r C_{v,i}^r = c_v^r \quad & \forall v \in N \\
\sum_i \lambda_{v,i}^{r+1} C_{v,i}^{r+1} = c_v^{r+1} \quad & \forall v \in N \\
\sum_i \Lambda_{i,j}^v = \lambda_{v,j}^{r+1} \quad & \forall v \in N, j \in [m_{r+1}] \\
\sum_j \Lambda_{i,j}^v = \lambda_{v,i}^r \quad & \forall v \in N, i \in [m_r] \\
\lambda_{v,i}^r, \lambda_{v,j}^{r+1} \in \text{SOS2} \quad & \forall v \in N, i \in [m_r], j \in [m_{r+1}] \\
\lambda_{v,i}^r, \lambda_{v,j}^{r+1}, \Lambda_{i,j}^v \in [0, 1] \quad & \forall v \in N, i \in [m_r], j \in [m_{r+1}] \\
\sum_i \lambda_{v,i}^r = \sum_i \lambda_{v,i}^{r+1} = 1 \quad & \forall v \in N
\end{aligned} \tag{\mathcal{P}_2}$$

5.1 Evaluation

Using the predictions made on the QEPA and MFPA datasets we generated patrols for each of the patrol posts in both national parks. Samples of these can be seen in Figure 4 where we show a heat map of the patrol effort corresponding to the distribution over computed patrols around posts for both the protected areas. These are currently being evaluated for real world deployment in both QEPA an MFPA. To evaluate the piecewise linear approximation with iWare-E prediction model, we look at the expected total predicted detections of illegal activity of the patrol schedules generated by the MILP. Given an optimal solution we can compute the actual predicted number of detected attacks using the iWare-E model. We then compare this prediction to the optimal objective value of the MILP used to compute c . These results are shown in Table 6 under the error column, where we measure the percent difference in these two values, averaged over all posts in the protected area. We see that we can get low approximation error when using the piecewise linear objective.

We also show the importance of being able to reason about continuous levels of patrol effort, where in Figure 5 we show the significant improvement in utility of the patrols computed, measured in terms of number of predicted detected attacks. For this comparison we let each breakpoint in the PWL approximation correspond to a discrete level of patrol effort and compared the number of predicted detections of both solutions. We see that even as we increase the number of levels of patrol effort to 80 levels, we still outperform the previous state-of-the-art by approximately 130% for the QEPA dataset and 150% for the MFPA dataset. Additionally the PWL objective allows us to be much more scalable; as an example, the previous state-of-the-art method requires 80 levels of discretization to achieve the same average utility (in terms of predicted detections) as 10 levels for QEPA and 5 for MFPA. This difference in discretization results in 400 \times , and 140 \times decrease in runtime for QEPA and MFPA respectively when using our method. We show similar improvements in runtime for more of these fixed utility comparison points in Figure 6 where it can be seen that it takes significantly more computational power for the previous state-of-the-art to match our results.

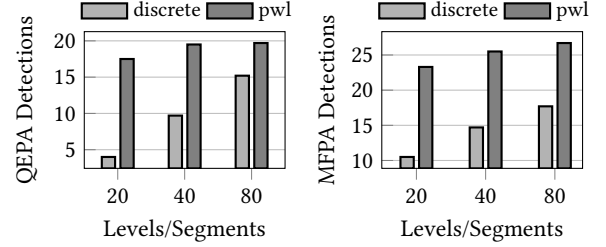


Figure 5: Improvement in solution quality of patrols planned for the QEPA and MFPA, using MILP (\mathcal{P}_2) compared to previous work using discrete levels of patrol effort. The utility is measured in number of predicted detected attacks.

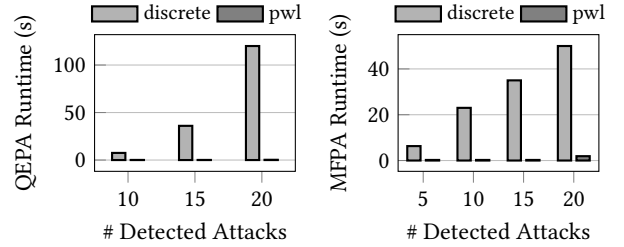


Figure 6: Improvement in runtime for computing patrols for QEPA and MFPA, using MILP (\mathcal{P}_2) compared to previous work using discrete levels of patrol effort for comparable solution quality, measured in predicted number of detected attacks. Results are averaged over 20 trials.

6 CONCLUSION

To make an impact in wildlife protection, it is crucial to adopt predictive and prescriptive models in the real fields. Previous works suffer from addressing the major technical and application challenges in this domain. In this paper, we present iWare-E, an efficient predictive model for wildlife protection, which accounts for imperfect crime information and uncertainty in wildlife data. This is the first time that this substantial challenge is addressed in data-driven adversarial reasoning in AI literature. Furthermore, we presented less computationally expensive fine-tuned generation of patrol routes based on the predictions of iWare-E to counteract poachers in the real-world more effectively (150% improvement in solution quality and 400 times higher speed). From domain perspective, previous works consider the coarse-grained temporal analysis of crime observations and they only evaluate on a single protected area. However, the predictive framework proposed in this paper significantly improves accuracy and runtime even for fine-grained analysis of crime over multiple protected areas. To our knowledge, this is the first adversary behavior model for wildlife protection that has been developed and evaluated at this scale in two protected areas to prove country-wide reliability in prediction results. Such predictive and prescriptive analysis can be invaluable for assisting law enforcement agencies in protecting wildlife more intelligently. **Acknowledgment:** This research was supported by MURI W911NF-11-1-0332 and NSF grant with Cornell University 72954-10598. We are grateful to the Wildlife Conservation Society and the Uganda Wildlife Authority for supporting data collection in Uganda.

REFERENCES

- [1] Nicola Basilico and Nicola Gatti. 2014. Strategic guard placement for optimal response to alarms in security games. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*. International Foundation for Autonomous Agents and Multiagent Systems, 1481–1482.
- [2] Nicola Basilico, Nicola Gatti, and Francesco Amigoni. 2009. Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *AAMAS*.
- [3] Michael J Chase, Scott Schlossberg, Curtice R Griffin, Philippe JC Bouché, Sinyehu W Djene, Paul W Elkan, Sam Ferreira, Falk Grossman, Edward Mtarima Kohi, Kelly Landen, et al. 2016. Continent-wide survey reveals massive decline in African savannah elephants. *PeerJ* 4 (2016), e2354.
- [4] Yuehui Chen, Bo Yang, and Ajith Abraham. 2007. Flexible neural trees ensemble for stock index modeling. *Neurocomputing* 70, 4 (2007), 697–703.
- [5] Rosie Cooney, Dilys Roe, Holly Dublin, Jacob Phelps, David Wilkie, Aidan Keane, Henry Travers, Diane Skinner, Daniel WS Challender, James R Allan, et al. 2017. From poachers to protectors: engaging local communities in solutions to illegal wildlife trade. *Conservation Letters* 10, 3 (2017), 367–374.
- [6] R Critchlow, AJ Plumpton, M Driciru, A Rwetsiba, EJ Stokes, C Tumwesigye, F Wanyama, and CM Beale. 2015. Spatiotemporal trends of illegal activities from ranger-collected data in a Ugandan national park. *Conservation Biology* 29, 5 (2015), 1458–1470.
- [7] Fei Fang, Thanh H Nguyen, Rob Pickles, Wai Y Lam, Gopalasamy R Clements, Bo An, Amandeep Singh, Milind Tambe, and Andrew Lemieux. 2016. Deploying PAWS: Field optimization of the protection assistant for wildlife security. In *IAAI*.
- [8] Fei Fang, Peter Stone, and Milind Tambe. 2015. When security games go green: Designing defender strategies to prevent poaching and illegal fishing. In *IJCAI*.
- [9] Shahrzad Gholami, Benjamin Ford, Fei Fang, Andrew Plumpton, Milind Tambe, Margaret Driciru, Fred Wanyama, Aggrey Rwetsiba, Mustapha Nsubaga, and Joshua Mabonga. 2017. Taking it for a test drive: a hybrid spatio-temporal model for wildlife poaching prediction evaluated through a controlled field test. In *Proceedings of the European Conference on Machine Learning & Principles and Practice of Knowledge Discovery in Databases, ECML PKDD*.
- [10] Shahrzad Gholami, Bryan Wilder, Matthew Brown, Arunesh Sinha, Nicole Sintov, and Milind Tambe. 2016. A game theoretic approach on addressing cooperation among human adversaries. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems*.
- [11] Shahrzad Gholami, Bryan Wilder, Matthew Brown, Dana Thomas, Nicole Sintov, and Milind Tambe. 2016. Divide to Defend: Collusive Security Games. In *GameSec*. Springer, 272–293.
- [12] Guo Haixiang, Li Yijing, Jennifer Shang, Gu Mingyun, Huang Yuan Yue, and Gong Bing. 2017. Learning from class-imbalanced data: review of methods and applications. *Expert Systems with Applications* 73 (2017), 220–239.
- [13] Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. 1991. Adaptive mixtures of local experts. *Neural computation* 3, 1 (1991), 79–87.
- [14] Debarun Kar, Fei Fang, Francesco Delle Fave, Nicole Sintov, and Milind Tambe. 2015. A game of thrones: when human behavior models compete in repeated Stackelberg security games. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 1381–1390.
- [15] Debarun Kar, Fei Fang, Francesco Delle Fave, Nicole Sintov, and Milind Tambe. 2015. "A Game of Thrones": When Human Behavior Models Compete in Repeated Stackelberg Security Games. In *AAMAS*.
- [16] Debarun Kar, Benjamin Ford, Shahrzad Gholami, Fei Fang, Andrew Plumpton, Milind Tambe, Margaret Driciru, Fred Wanyama, Aggrey Rwetsiba, Mustapha Nsubaga, et al. 2017. Cloudy with a Chance of Poaching: Adversary Behavior Modeling and Forecasting with Real-World Poaching Data. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 159–167.
- [17] Christopher Kiekintveld, Towhidul Islam, and Vladik Kreinovich. 2013. Security games with interval uncertainty. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*. International Foundation for Autonomous Agents and Multiagent Systems, 231–238.
- [18] Dmytro Korzhuk, Vincent Conitzer, and Ronald Parr. 2011. Solving Stackelberg games with uncertain observability. In *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 3*. International Foundation for Autonomous Agents and Multiagent Systems, 1013–1020.
- [19] Wee Sun Lee and Bing Liu. 2003. Learning with positive and unlabeled examples using weighted logistic regression. In *ICML*, Vol. 3.
- [20] Berendien Anna Lubbe, Elizabeth Ann du Preez, Anneli Douglas, and Felicity Fairer-Wessels. 2017. The impact of rhino poaching on tourist experiences and future visitation to National Parks in South Africa. *Current Issues in Tourism* (2017), 1–8.
- [21] Sara McCarthy, Milind Tambe, Christopher Kiekintveld, Meredith L. Gore, and Alex Killion. 2016. Preventing Illegal Logging: Simultaneous Optimization of Resource Teams and Tactics for Security. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI'16)*. AAAI Press, 3880–3886. <http://dl.acm.org/citation.cfm?id=3016387.3016450>
- [22] Richard McKelvey and Thomas R. D.Palfrey. 1995. Quantal Response Equilibria for Normal Form Games. In *Games and Economic Behavior*.
- [23] Enrique Munoz de Cote, Ruben Stranders, Nicola Basilico, Nicola Gatti, and Nick Jennings. 2013. Introducing alarms in adversarial patrolling games. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*. International Foundation for Autonomous Agents and Multiagent Systems, 1275–1276.
- [24] Robin Naidoo, Brendan Fisher, Andrea Manica, and Andrew Balmford. 2016. Estimating economic losses to tourism in Africa from the illegal killing of elephants. *Nature communications* 7 (2016).
- [25] Thanh H Nguyen, Francesco M Delle Fave, Debarun Kar, Aravind S Lakshminarayanan, Amulya Yadav, Milind Tambe, Noa Agmon, Andrew J Plumpton, Margaret Driciru, Fred Wanyama, et al. 2015. Making the most of our regrets: Regret-based solutions to handle payoff uncertainty and elicitation in green security games. In *GameSec*. Springer, 170–191.
- [26] Thanh H Nguyen, Arunesh Sinha, Shahrzad Gholami, Andrew Plumpton, Lucas Joppa, Milind Tambe, Margaret Driciru, Fred Wanyama, Aggrey Rwetsiba, Rob Critchlow, et al. 2016. CAPTURE: A new predictive anti-poaching tool for wildlife protection. *AAMAS*, 767–775.
- [27] Thanh Hong Nguyen, Rong Yang, Amos Azaria, Sarit Kraus, and Milind Tambe. 2013. Analyzing the Effectiveness of Adversary Modeling in Security Games. In *AAAI*.
- [28] Steven Okamoto, Noam Hazan, and Katia Sycara. 2012. Solving non-zero sum multiagent network flow security games with attack costs. In *AAMAS*, 879–888.
- [29] Nazneen Fatema Rajani and Raymond J Mooney. 2016. Supervised and Unsupervised Ensembling for Knowledge Base Population. *arXiv preprint arXiv:1604.04802* (2016).
- [30] Wei Wei, Jinju Li, Longbing Cao, Yuming Ou, and Jiahang Chen. 2013. Effective detection of sophisticated online banking fraud on extremely imbalanced data. *World Wide Web* 16, 4 (2013), 449–475.
- [31] George Wittemyer, Joseph M Northrup, Julian Blanc, Iain Douglas-Hamilton, Patrick Omondi, and Kenneth P Burnham. 2014. Illegal killing for ivory drives global decline in African elephants. *Proceedings of the National Academy of Sciences* 111, 36 (2014), 13117–13121.
- [32] Haifeng Xu, Benjamin Ford, Fei Fang, Bistra Dilkina, Andrew Plumpton, Milind Tambe, Margaret Driciru, Fred Wanyama, Aggrey Rwetsiba, Mustapha Nsubaga, et al. 2017. Optimal Patrol Planning for Green Security Games with Black-Box Attackers. In *International Conference on Decision and Game Theory for Security*. Springer, 458–477.
- [33] Rong Yang, Benjamin Ford, Milind Tambe, and Andrew Lemieux. 2014. Adaptive resource allocation for wildlife protection against illegal poachers. In *AAMAS*.
- [34] Rong Yang, Christopher Kiekintveld, Fernando Ordonez, Milind Tambe, and Richard John. 2011. Improving resource allocation strategy against human adversaries in security games. In *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*, Vol. 22. 458.