Data-Driven Multimodal Patrol Planning for Anti-poaching

Weizhe Chen[1], Weinan Zhang[1], Duo Liu[2], Weiping Li[2], Xiaojun Shi[3], Fei Fang[4]
[1] Shanghai Jiao Tong University,
[2] WWF-China NEC,
[3] Huang Ni He Forest Bureau,
[4] Carnegie Mellon University
chenweizhe@sjtu.edu.cn, wnzhang@apex.sjtu.edu.cn, {dliu, wpli}@wwfchina.org, 15043342333@139.com,
feif@cs.cmu.edu

Abstract

Wildlife poaching is threatening key species that play important roles in the ecosystem. With historical ranger patrol records, it is possible to provide data-driven predictions of poaching threats and plan patrols to combat poaching. However, the patrollers often patrol in a multimodal way, which combines driving and walking. It is a tedious task for the domain experts to manually plan such a patrol and as a result, the planned patrol routes are often far from optimal. In this paper, we propose a data-driven approach for multimodal patrol planning. We first use machine learning models to predict the poaching threats and then use a novel mixed-integer linear programming-based algorithm to plan the patrol route. In a field test focusing on the machine learning prediction result at Jilin Huangnihe National Nature Reserve (HNHR) in December 2019, the rangers found 42 snares, which is significantly higher than the historical record. Our offline experiments show that the resulting multimodal patrol routes can improve the efficiency of patrol and thus they can serve as the basis for future deployment in the field.

Introduction

Poaching has been a global concern for many years. Key species that play important roles in the ecosystem, such as elephants, tigers, rhinos, are under constant poaching threats due to the trade of profitable wildlife products, including ivory, rhino horns, and tiger skins (Chase et al. 2016; Spillane 2016). Besides, wild boars, roar deer, and other wildlife species are poached for bushmeat (Warchol 2004). Wildlife conservation agencies are trying to protect wildlife from poaching by sending rangers to patrol in the conservation sites to confiscate poaching tools like snares and stop the poachers (Lemieux 2014).

The patrol routes are traditionally planned by domain experts such as site managers and experienced rangers, with the main goal of detecting and deterring poaching activities. In many huge conservation sites, the rangers patrol through multiple modes of transportation, e.g., driving through cars and motorcycles, walking, and sometimes even through boats. We consider a typical multimodal patrol scenario: rangers drive along existing roads to open trails to a drop-off location, then patrol on foot and get back to the vehicle at a pick-up location, and then drive to the next drop-off location, etc. If the ranger team has a driver, the pick-up location can be different from where the rangers are dropped off. Such a complex patrol makes route planning a tough task as the planner needs to take into account the varying poaching threat across the site, the movement speed and the total patrol time. In current practice, the domain experts rely on their understanding of the poaching threat and some rudimentary planning methods to plan the routes, resulting in a significant cognitive burden on the domain experts. Furthermore, their understanding may be biased and the planned routes may be far from optimal. Besides, when there is employee movement, it is hard to transfer the site-specific knowledge needed for planning the routes to new managers and rangers who lack experience patrolling in the specific site.

In this paper, we propose a data-driven approach to plan the patrol routes, aiming to make better use of the limited patrol time of rangers and reduce the cognitive burden of patrol planning. In some sites, rangers record their waypoints and findings during the patrols through the Spatial Monitoring And Reporting Tool (SMART), including where and when they find animal signs, human activity signs, and snares. These records enable the use of machine learning (ML) methods to analyze the data. Some prior efforts use ML to assist conservation agencies identify poaching hotspots (Nguyen et al. 2016; Gholami et al. 2017, 2020), but the results are still far from perfect. Also, most of them focus on predicting the poaching threat and do not take planning into account. Among the ones that consider route planning, the practical aspect of multimodal patrol is neglected. In contrast, in our approach, we first apply ML methods to the processed data to predict poaching threats and then use mathematical programming to plan multimodal patrols. The workflow is shown in Figure 1. In the poaching threat prediction part, we use a averaging ensemble model with a neural network (NN) and decision trees to predict the poaching threats. We show that with proper hyper-parameter tuning, NN can improve the F1 score by 0.2 in our dataset of Jilin Huangnihe National Nature Reserve (HNHR), a conservation area of about 75 sq. km in Northeast China. Furthermore, the averaging ensemble model can lead to a precision of 0.63. During a field test at HNHR focusing on the prediction, the rangers found 42 snares placed by poachers, encountered other humans (who are likely poachers) 3 times,

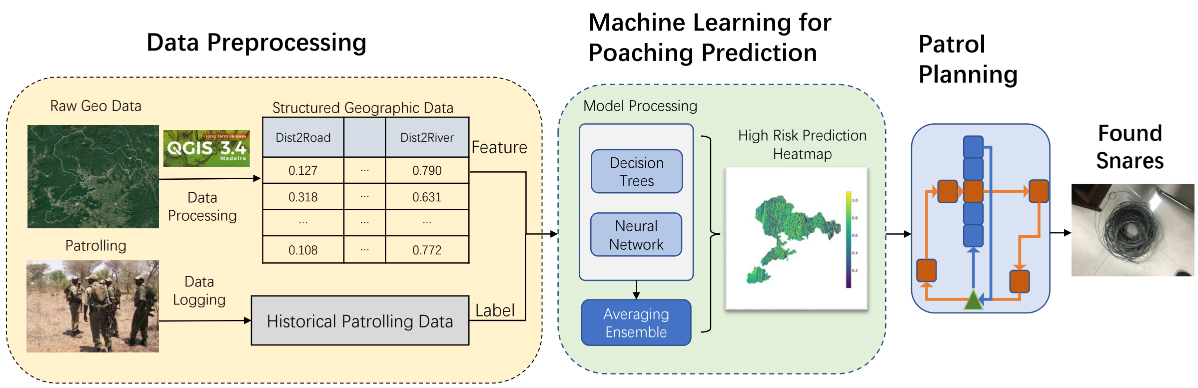

Figure 1: The workflow of this study. The images in this figure are just for example and not our result.

and found 9 other human activity signs. These results are significantly better than what has been shown in previous work. In the patrol planning part, we use a novel mixed-integer linear programming (MILP)-based algorithm to optimize multimodal patrol routes. The experimental results show that the routes found using this approach can improve patrol efficiency. We believe this work can serve as the basis for further deployment.

Related Work

There have been several lines of work on using artificial intelligence techniques for anti-poaching. Some of them focus on building game-theoretic models to analyze the strategic interaction between poachers and conservation agencies. In (Fang et al. 2015), the green security game model was proposed as an extension of the Stackelberg Security Game (Tambe 2011), which assumes the poachers respond to a linear combination of past patrol and current patrol, and plans patrols with limited lookahead. Building upon this model, Wang et al. (2019) further took into consideration the real-time information , and Huang et al. (2020) introduced the aspect of community engagement. Recently, Gholami et al. (2019) studied patrol planning in an online setting. Unfortunately, most of these works cannot be deployed in the real world. An important reason is that the strong assumptions in these models may not hold in practice. For example, the poachers are assumed to respond to rangers' patrol perfectly rationally or following the quantal-response model (McKelvey and Palfrey 1995) or its variants. However, the poachers often cannot perfectly observe the patrol strategies used by the rangers and can hardly make (boundedly) rational decisions based on them. Another reason is the computational intractability of game-theoretic approaches when applied to large parks.

Another line of work used ML to predict poaching activities, which have led to field tests and field deployments with promising results in multiple countries. INTERCEPT (Kar et al. 2017) used an ensemble of decision trees and achieved a better result than the previous model CAPTURE, which used dynamic Bayesian networks (Nguyen et al. 2016). It was tested in Queen Elizabeth National Park (QENP) with one poached elephant and three sets of snares found. Gholami et al. (2017) introduced a geo-clustering technique and produced a hybrid model consisting of Markov random fields and decision trees with a bagging ensemble. Recently, a new ensemble method iWare-E (Gholami et al. 2018) was proposed, using other ML models as weak learners. Gholami et al. (2020) further introduced Gaussian process (GP) (Rasmussen 2003) to provide predictions as well as confidence intervals. Most of the work in this line studied how to predict the poaching threat and leave the planning part to domain experts. Our work is closely connected to (Gurumurthy et al. 2018) as they also use data from HNHR, but our work is different from it in several aspects. Firstly, we focus on planning routes based on the prediction results while (Gurumurthy et al. 2018) only considers predicting the poaching threats. Secondly, our dataset has more features, including ones that we extract from satellite imagery. Lastly, we improve the machine learning methods for prediction by using an averaging ensemble and tune the hyperparameters carefully.

There are a few existing works that employ machine learning to predict poaching hotspots then plan patrol (Fang et al. 2015; Gholami et al. 2018, 2019; Xu et al. 2017). However, they only considered patrolling on foot and did not take into account the multimodal patrol problem. In this paper, we follow this predict-and-plan framework and focus on multimodal patrol planning to cover more areas with high poaching threats. It is a very important future direction to extend our work to consider the long-term impact of patrolling on poaching.

Research in multimodal transportation is also related to our work. These works investigate how to combine multiple transporting modes, including rails, air, water modes, to send cargo from points to points (SteadieSeifi et al. 2014). (Alumur, Yaman, and Kara 2012) provided a MILP-based algorithm to determine the location of the hub facilities. However, many existing works assume a hub-and-spoke structure and focus on planning a path from origin to destination to minimize travel time or cost. In contrast, in our work, the rangers can walk freely in the protected area, thus presenting no such special structure to be exploited. In addition, the goal of patrols in our problem is not to go from an origin to a destination but to traverse through areas with high poaching risk to detect poaching activities. These differences require a completely new formulation.

Data Preparation

To construct the dataset for ML-based prediction, we first discretize the geographical area into a grid, with grid cell size $200m \times 200m$. Then, we extract the patrolling record from SMART, which contains waypoints and detailed information such as time, coordinates, and indicators of whether there are any poaching signs (usually snares) found. In our constructed dataset, each data point corresponds to one grid cell in one year. A data point is labeled negative if rangers patrolled in that cell that year but did not find any poaching activities, and positive if there are poaching activities recorded.

We further extract geospatial features. In the case study in HNHR, we used the following features: (i) the closest distance to stream, ridgeline, river, marsh, different types of land (naming landcover$_{1-10}$); (ii) the closest distance to the nearby village, farmland, village road, provincial road, national road, highway, and the boundary of the conservation site; (iii) elevation and slope. It is a significant amount of work to collect the related information from local agencies to extract these features, and some of the information is not available in digital form (e.g., the shape of the farmland). To reduce the workload of site managers and local agencies, and get a richer set of features, we extract the distance-related features from satellite imagery in a semi-automated way. We first use the ROI function in Exelis Visual Information Solutions (ENVI) to extract the shapefiles. Then, we obtain the distance information using functions in Geographical Information System (GIS) software products, e.g., the NNJoin function in QGIS. The elevation-related features are extracted from the digital elevation model.

In patrol planning, we use the distance to road features to specify where driving is feasible.

Machine Learning for Poaching Prediction

In this section, we describe the machine learning model we use for poaching threat prediction.

Data Augmentation

While datasets in the anti-poaching domain are often constructed similarly, they share some common challenges. First, as the rangers cannot perfectly detect all poaching activities, a negative label does not mean there is no poaching activity in the corresponding grid cell. Also, as the snares found are very sparse in the whole studied area, the datasets suffer from significant label imbalance. In our dataset, the number of positive data points is about 1% of that of the negative data points. To overcome these challenges, we use a few data augmentation methods.

Positive Data Upsampling. Since the number of positive data is very small, we duplicated the positive data several times in training. With this approach, we ensure the number of positive examples and the negative examples remains at the same level.

Positive Labeling via Domain Experts. Domain experts such as conservation site managers often know which areas have more poaching threats based on their conversations with former poachers, local villagers etc. While such understanding can be biased, it complements the recorded patrol data. So we augment the dataset based on domain experts knowledge. In a previous study (Gurumurthy et al. 2018), domain experts provided scores for different areas indicating their estimation of the poaching threat. Using these experts' provided scores, we choose the data points with the highest average scores and label them as positive to double the number of positive data points in the dataset.

Ensemble Learning

In practice, rangers care more about precision when the patrol resource is limited, as they can only patrol very few areas. To improve the precision, we use an ensemble consisting of three base machine learning models: the first is a neural network-based model, the second is a bagging decision tree (Breiman 1996), and the third is a gradient boosting decision tree predictor implemented with XGBoost (Chen and Guestrin 2016). We choose these three base models for their computational efficiency and performance on the HNHR site, as detailed later. We use simple averaging to compute the final prediction of the ensemble after thresholding, i.e., the final prediction is proportional to the number of base models whose predictions are above the model-specific thresholds. The normalize the final prediction value to $[0, 1]$. It is possible to use other base models for the ensemble as well.

Multimodal Patrol Planning

In this section, we detail our MILP-based approach to plan the multimodal patrol. We use a coarser grid with cell size $1km \times 1km$ in patrol planning for a few reasons. First, a coarser grid gives the rangers more flexibility to the rangers. Such flexibility makes it easier for them to adopt the recommended patrol routes and can motivate them to use their domain knowledge to find the poaching signs. For example, the rangers may know that it is easier to find snares near certain types of trees, and they can look for these trees during patrols. Second, it is computationally challenging to find the optimal patrol route with a fine grid, as the number of possible patrol routes increases exponentially with the number of grid cells. We refer to a grid cell in this coarse grid as a block. For block i in this coarse grid, we use p_i to indicate the predicted poaching threat level. p_i is calculated as the average value of the ML-model's prediction among the 25 small grid cells of size $200m \times 200m$ within this block i. We check if the smallest value of the Dist2Road feature in the small grid cells is less than 0.1 to determine if driving is feasible in this block. We denote the set of blocks where driving is feasible as Φ.

The MILP is based on a segment-unrolled graph, which leverages insights from the time-unrolled graphs used in existing patrol route planning work (Xu et al. 2017). We observe that any patrol route can be described as a sequence of route segments where each segment represents the patroller moving from one block to another either by car or on foot, i.e., the mode of transportation cannot be changed within a segment. But a route segment can connect two blocks i, j that are far away as long as it is possible to move from i to j using a single mode of transportation. The ranger can patrol in a block j after completing a route segment from block i to j. The ranger may also complete two consecutive route segments from i to j and from j to k without patrolling in block j. Thus, we construct the segment-unrolled graph G which has $N \times (K+1)$ nodes where N is the number of blocks and K is the maximum number of route segments in a feasible patrol route. Node $v_{i,k}$ in G denotes the block i that serves as the starting block of the k^{th} route segment. The edges in G only exist between nodes corresponding to two consecutive segments. Edge $e = (v_{i,k}, v_{j,k+1})$ indicates that a possible k^{th} route segment is to move from block i to block j. Since it is always possible to move from a block i to any other block j on foot, there exists an edge $e = (v_{i,k}, v_{j,k+1})$ for any $i, j \in \{1..N\}, k \in \{1..K\}$. In addition, if it is possible to drive from block i to j, we add a special attribute to all the edges of $e = (v_{i,k}, v_{j,k+1}), k = 1..K$ to indicate that these route segments can be completed through driving. To find these driving-compatible route segments, we first notice that the starting block of a driving-compatible edge has to be in Φ, and driving is feasible in these blocks. Also, the patroller can switch from driving mode to walking mode or vice versa in these blocks. We then apply the Floyd algorithm to check if block i and j are connected through roads.

An example is shown in Figure 2. In this example, there are three blocks in the area of interest. Block 1 and 2 are connected through roads. If the patrol route can have at most 4 segments, we construct the graph G with 15 nodes. The dashed edges indicate driving-compatible route segments. For example, the dashed edge between node $(1, 1)$ and $(2, 2)$ indicate that the first route segment of the patrol route can be moving from block 1 to block 2 through driving.

In practice, the constraints on the patrol route are usually not on the number of route segments. We consider the following practical constraints: Assuming the time needed to patrol in one block is 1, the total time for completing the patrol route has to be less than T. The time spent walking during the patrol has to be less than L_x. The patrol route starts from and ends with the block s where the patrol post locates. In addition, a subset of blocks $M \subset \{1..N\}$ have to be visited but not necessarily patrolled. M often represents the blocks where the camera traps are placed and the rangers need to replace film or battery. Given these constraints, we set the value of K to be $3T$, which is an upper bound of the number of route segments in a patrol route as after patrolling in one block, the patrollers may need to take at most three route segments in walking, driving and talking mode to get to another block to patrol.

Now we introduce some of the notations used in the MILP formulation. We use the walking distance between two neighboring blocks as the unit distance. l_e where $e = (v_{i,k}, v_{j,k+1})$ denotes the walking distance between block i to block j calculated using Manhattan distance. l_e^r denotes

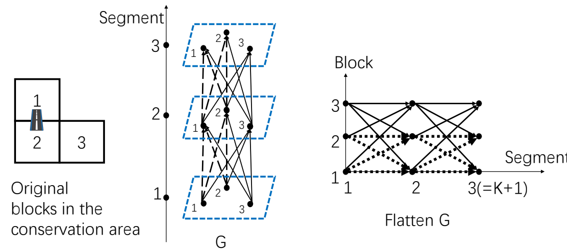

Figure 2: An example of graph G. There is a road between block 1 and block 2. The dotted lines in G means both types of edge exist between nodes, and solid lines indicates only the edge of walking exists between nodes.

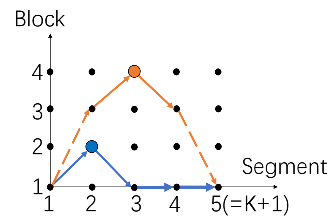

Figure 3: An example of two possible routes on G. Orange and blue arrows show two routes, where dotted arrows indicate driving and solid arrows indicate walking. Colored nodes means that the ranger patrols in the corresponding block.

the driving distance between i and j which is pre-computed with Floyd algorithm. Clearly $l_e^r \geq l_e$ if l_e^r exists as one can always walk along the road. $\sigma^+(v_{i,k})$ and $\sigma^-(v_{i,k})$ denote the sets of edges that enter into or exit from node $v_{i,k}$ respectively. α and β represent the time needed for walking and driving for one unit distance.

The time needed to complete the first k segments in two patrol routes may not be the same as the distance, the transportation mode chosen, and the number of blocked patrolled can be different. Take the two routes in Fig. 3 as an example. Assume the walking distance between the first and the second block is 1, the driving distance between the first and the third block is 5, and the walking distance between the third block and the fourth block is 2. And assume $\alpha = 0.01, \beta = 0.1$. For the orange route, the time needed for completing the first two segments is $\alpha \times 5 + \beta \times 2 = 0.25$. And for the blue route, the time needed is $\beta \times (1 + 1) + 1 = 1.2$ as the rangers spend one unit of time to patrol in block 2.

Next, we define the variables used in the MILP. Let $f_{e,k} \in \{0, 1\}$ and $f_{e,k}^r \in \{0, 1\}$ denote whether the k^{th} route segment is taking edge e in walking and driving mode respectively. $x_{i,k} \in \{0, 1\}$ denotes whether the rangers are at block i after completing the $(k-1)^{th}$ route segment and are in walking mode in the k^{th} segment. Similarly, $y_{i,k} \in \{0, 1\}$ denotes whether the rangers are at block i after the $(k-1)^{th}$ segment and are in driving mode in the k^{th}segment. $z_{i,k} \in \{0, 1\}$ denotes whether the rangers spend time to patrol in block i after completing the $(k-1)^{th}$ route segment. If $z_{i,k} = 0$ and $x_{i,k} + y_{i,k} = 1$, the rangers only traverse through block i without patrolling in it.

Given these variables, the MILP formulation is as follows:

$$\max_{x,y,f,f^r,z} \sum_{i=1}^{N}\sum_{k=1}^{K} p_i \times z_{i,k} \tag{1}$$

$$\text{s.t.} \sum_{i=1}^{N}\sum_{k=1}^{K} x_{i,k} \le L_x, \tag{2}$$

$$\sum_{k=1}^{K}\sum_{i=1}^{N} z_{i,k} + \sum_{e}(\alpha l_e^r f_{e,k}^r + \beta l_e f_{e,k}) \le T, \tag{3}$$

$$\forall i,k, x_{i,k} = \sum_{e\in\sigma^+(v_{i,k})} f_{e,k} \tag{4}$$

$$\forall i,k, y_{i,k} = \sum_{e\in\sigma^+(v_{i,k})} f_{e,k}^r, \tag{5}$$

$$\forall i, \sum_{k=1}^{K} z_{i,k} \le 1, \tag{6}$$

$$\forall i \in M, \sum_{k=1}^{K} z_{i,k} \ge 1, \tag{7}$$

$$\sum_{e\in\sigma^-(s)} f_{e,0}^r = \sum_{e\in\sigma^+(s)} f_{e,K}^r = 1, \tag{8}$$

$$\begin{aligned}\forall i \in \Phi, \forall k, &\sum_{e\in\sigma^+(v_{i,k})} f_{e,k} + f_{e,k}^r \\ &= \sum_{e\in\sigma^-(v_{i,k})} f_{e,k} + f_{e,k}^r,\end{aligned} \tag{9}$$

$$\forall i \notin \Phi, \forall k, \sum_{e\in\sigma^+(v_{i,k})} f_{e,k} = \sum_{e\in\sigma^-(v_{i,k})} f_{e,k}, \tag{10}$$

$$\forall k, \sum_{i=1}^{N} x_{i,k} + y_{i,k} \le 1 \tag{11}$$

$$\forall i,k, z_{i,k} \in \{0,1\}, x_{i,k} \ge 0, y_{i,k} \ge 0, x_{i,k} \ge z_{i,k}, \tag{12}$$

$$\forall e,k, f_{e,k} \ge 0, f_{e,k}^r \ge 0 \tag{13}$$

The objective is to maximize the sum of the predicted poaching threat level p_i in all patrolled blocks. Constraint 2 limit the total time to patrol. Constraint 3 limits the total time for completing the patrol route, where the first term is the time for patrolling, and the second term is the time for walking and driving. Constraint 4 and 5 ensures the rangers' location as indicated by variables x and y is consistent with the patrol route as indicated by variables f and f^r. Constraint 6 ensures that any block is patrolled at most once. Constraint 7 ensures the must-visit points are patrolled. Constraint 8 limits the starting point and the endpoint of the patrol route to be the block where the patrol post locates. Constraint 9 and 10 ensure the route segments are connected. Constraint 11 is used to ensure the rangers take the k^{th} route segment in either walking or driving mode.

	Precision	Recall	F1	L&L	AUC
NN	0.44	**0.46**	**0.44**	15.5	0.83
DT	0.61	0.33	0.43	**15.8**	0.80
XGBoost	0.58	0.30	0.39	13.1	0.82
Ensemble	**0.63**	0.29	0.40	14.3	0.72
Bagging GP	0.14	0.78	0.24	9.18	**0.86**
DT_{ref}	0.18	0.31	0.23	15.4	–
NN_{ref}	0.02	0.79	0.04	3.70	–
NN_{small}	0.17	0.25	0.20	3.40	0.80
$\text{NN}_{\text{less epoch}}$	0.33	0.37	0.35	9.65	0.83
$\text{XGBoost}_{\text{alter}}$	0.28	0.33	0.30	7.35	0.77

Table 1: The offline results, where the ref result is from (Gurumurthy et al. 2018) which use a slightly different dataset.

#Models as Positive	0	1	2	3
Count	52784	1521	867	1762

Table 2: The statistics of predictions of cells.

Case Study on HNHR

Offline Experiment for Prediction

Dataset In this paper, we use the dataset of HNHR. The raw data contains the patrol information of the last 5 years (2014 - 2018) in SMART formats. All features are linearly normalized to $[0, 1]$ for better computational results.

To test our models, we build a test set by randomly choosing 25% of the grid cells and use the labeled data corresponding to those cells in 2018. The data corresponding to the other 75% of the grid cells in 2015-2017 is used as the training set. We use such a split since it may be easy for a model to predict the high threat at a cell if poaching activities were found in that exact cell in other years. We want to test if our model is able to predict poaching activities in cells that do not show up in the training data. Also, we split the temporal dimension to test if our model can predict the future properly. To tune the hyper-parameters, we separately split the positive data and the negative data into 4 folds, and 1 fold each is chosen as the validation set.

Metrics Following the literature, we evaluate our models by precision, recall, F1 score, L&L score (defined as Recall×TestSetSize/PredictedPositives), and area under ROC curve (AUC) score. These metrics are common in the anti-poaching field because of their strong link to the real world. Among these metrics, we found that only using the AUC will not be very informative due to the data imbalance. However, because the first 4 metrics are threshold-sensitive, and therefore time-consuming, we first use AUC score of validation data to roughly estimate the performance of our models and further enumerate the threshold to select the best one and compare between different models.

Implementation Details *Data Augmentation* We add precisely 80 positive data points using the positive labeling via domain experts, and we duplicate the positive data 175 times to upsample our positive data.

Neural Network (NN). The neural network model contains 5 hidden fully-connected layers. It is trained for 400 epochs.

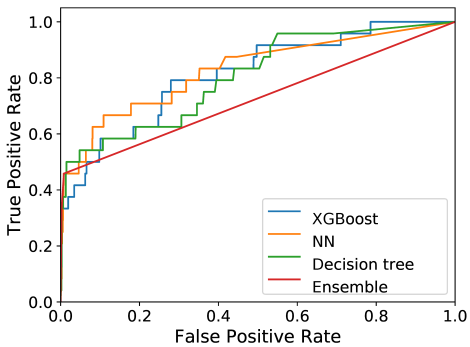

Figure 4: The ROC Curve of our models.

Bagging Decision Tree (DT). We use a bagging ensemble decision tree algorithm with 1,000 trees where each tree is trained using only 20% of the data.

Gradient Boosting Decision Tree (XGBoost). We specify the following parameters in training: 'max_depth': 20, 'eta': 0.3, 'silent': 1.

Averaging Ensemble (Ensemble). We use the threshold that resulted in the highest F1 score in the validation set of each model as the threshold for transforming the real-value prediction to a binary prediction. We choose an ensemble consisting of XGBoost, DT, and NN since it performs the best.

Results We report the performance of the base models we use and our ensemble model in the first part of Table 1. The four threshold-sensitive metrics are obtained by enumerating over the threshold and selecting the threshold that maximizes the F1 score. Not surprisingly, the ensemble method leads to higher precision compared to the base models. As shown in Table 2, about 5% data is predicted positive in one of the three models, and only about 2% of the data are predicted positive in all three models. This means the high-threat areas predicted by the base models partially overlap. When the patrol resources are extremely limited, it is reasonable to spend more effort on the overlapped predicted high-threat areas. The ROC curves of the base models and the ensemble method are shown in Figure 4.

We also compare our models with several baselines, shown in the second part of Table 1. The first baseline method is Bagging Gaussian Process (GP). GP can provide both predictions as well as confidence intervals, and it is shown to have decent performance in several other conservation sites (Gholami et al. 2020). However, due to the large number of grid cells in our case study, it is computationally intractable to train a single model with all the data. Therefore, we use a bagging method for GP. We train 20 models each with 10% of the training data (chosen randomly) and provides the final prediction by averaging the 20 models. We also compare the performance of our models with the results reported in Gurumurthy et al. (2018), denoted with subscription ref in the table. As we can see from these results, the performance of the neural network model after careful

(a) Found snares (b) Field deployment

Figure 5: (a) The snares found this year; (b) the visualization of a part of our field deployment, where the color in the rectangles is the heatmap of our prediction.

hyper-parameter tuning is much better than any previous result reported on a similar data set for the same conservation site (Gurumurthy et al. 2018). Surprisingly, GP does not work well in HNHR.

As an ablation study, we change the capacity of the neural network and the number of training epochs, with results shown in the third part of Table. 1. The naive neural networks ($\mathrm{NN}_{\mathrm{small}}$) are a 2-layer fully connected hidden layer , the less trained model ($\mathrm{NN}_{\mathrm{less\ epoch}}$) is only trained for 40 epochs, and the alternate XGBoost model ($\mathrm{XGBoost}_{\mathrm{alter}}$) has the parameters of : 'max_depth': 40, 'eta': 0.1, 'silent': 1.. We found that increasing the training epochs is indeed helpful. Also, using 5 hidden layers instead of 2 leads to better performance. XGBoost's performance also degrades when the parameters are not carefully chosen. We also test the result of not normalizing the data to $[0, 1]$, and the result is almost the same with Table. 1, with a difference less than 0.05 in all cells.

Field Deployment for Prediction

We use the ensemble model to predict the poaching threat in the whole area of HNHR and get a prediction for every single $200m \times 200m$ cell in the area. We visualize the result through heatmaps. Then we manually choose a few rectangular-shaped regions to patrol based on the predictions. Each selected region is $5km \times 1km$ in size, and the rangers can navigate in the region based on their own knowledge to find snares. We design the patrol regions considering both the high expectation value of finding snares and the need for exploration in areas that are rarely visited. These areas are tagged as high, medium, and low threat according to their mean value of prediction in the area. Detailed results are shown in Tab. 3.

Deployment Results The result was deployed in December 2019 for one week. In this one-week field test, the rangers found 42 snares in 4 such rectangle areas. Moreover, among those areas, there were a maximum 36 snares found in a single area, which is a high-threat area according to our model. They also bumped into a few people that are likely poachers. The snares found this year include new snares that are probably recently put by poachers and old snares that poachers may not be using. For comparison, Gurumurthy et al. (2018) conducted a field test for one month in the winter of 2017-2018 in the same conservation site with

Group type	#snares	Avg #snares	#human sign	Avg Prediction	Max Prediction	#Area
High Threat	41	3.416	1	[0.204, 1]	[0.66, 1]	12
Medium Threat	0	0	7	[0.0492, 0.204)	[0.33, 0.66]	13
Low Threat	1	0.076	4	[0, 0.0492)	[0, 0.33]	13
Total (new in 2019)	42	-	12	-	-	38
2017-2018 winter	29	-	4	-	-	-

Table 3: The result of our field deployment. Together With the prediction result of ensemble model in each single area.

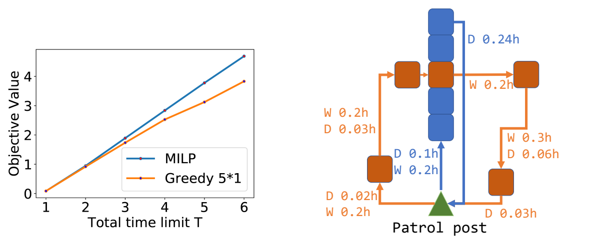

(a) The objective value.

(b) Illustration of the route.

Figure 6: The result of our route planning. (a) is the figure of expected threat value patrolled with different total time limit. (b) is an illustration of the route result. W denotes walking and D denotes driving.

promising results of 22 snares found, most of which are old snares. The results from the December 2019 deployment are even better in terms of the number of snares found and the number of encounters of humans. Detailed statistics of the deployment results are shown in Tab. 3. As a 5×1 rectangle-shaped area contains 125 $200m \times 200m$ grid cells, we compute the mean value and maximum value of the predicted threat levels of these 125 cells for each area and report the range of them for all the high, medium, low threat areas in the table. We can see that most of the snares and human activity signs are found in the high and medium threat areas. There are also some snares and human activity signs found in the low-threat area. As we take a closer look at these areas, we found that previous patrol has rarely covered them, making it very challenging for our model to make accurate predictions in these areas. We believe that incorporating exploration in the future design of patrol routes will be essential. These promising results make it possible to conduct more tests in the future, including the potential field tests for our multi-model patrol planning algorithm.

Offline Experiment for Planning

Although our MILP-based planning algorithm has not yet been deployed in the field, we run offline experiments in one region within HNHR. We set $\alpha = 0.01, \beta = 0.1$, which is calculated based on the assumption that it takes roughly 1 hour to patrol a block, and the driving speed is $100km/h$, and the walking speed is $10km/h$ when rangers are not patrolling. We vary the time limit T from 1 to 6, and set $L_x = L_y = T = \frac{K}{4}$. We relax the binary variables x and y to real-valued variables in implementation as it will not change the optimal value. The intuition is that if $0 < x_{i,k} < 1$ for some i, k, then $z_{i,k} = 0$ due to constraint 12. This means node p_i does not contribute to the objective value. So we can reduce $x_{i,k}$ and increase $x_{i',k}$ or $y_{i',k}$ for some i' without reducing the objective value and we can repeat the process until $x_{i,k} = 0$. We set 2 hours as the time limit for solving the MILP. This is a reasonable time limit the domain experts are willing to accept to get their patrol route. We compare it with a greedy method (referred to as 'Greedy 5*1'), which enumerates all possible $5km \times 1km$ rectangles and finds the one with the maximum total predicted poaching threat. This greedy method is also used to determine the high-threat areas in the aforementioned field deployment for prediction. The results are shown in Fig. 6. It can be seen that the MILP-based method can improve the efficiency of patrolling, and the improvement is more significant when we have a larger time limit of T.

For each instance, running the algorithm takes up to $60G$ memory. When it terminates at the 2-hour time limit, the optimality gap is in the range of $[0.01, 0.3]$. One important direction for future work is to improve the computational efficiency of the algorithm as the site managers often need to run the algorithm using a personal computer.

To show the planned route to the rangers and site managers, we automatically generate a heatmap of our prediction of poaching threat, overlaid with the suggested route. The rangers can download these results to their smartphones or print them out on paper so that they can easily follow the suggested routes during their patrols. We believe that this route planning method will serve as the basis for future deployment.

Conclusion

In this work, we provide a framework for data-driven multimodal patrol planning. In this framework, we first use an ensemble method with several base machine learning models to predict the poaching threat levels. The models are trained on historical patrol data with geospatial features extracted from satellite imagery. We then use a novel MILP-based algorithm to compute the patrol route with multiple transportation modes including walking and driving. In a case study of HNHR, we validate the poaching threat prediction model with a one-week field test. We further run offline experiments to show that the framework can help rangers find better patrol routes that cover more high-threat areas.

Acknowledgments

W. Zhang is supported by "New Generation of AI 2030" Major Project (2018AAA0100900) and National Natural Science Foundation of China (61702327, 61632017). F. Fang is supported in part by NSF grant IIS-1850477. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the funding agencies. Also, we thank rangers from Huang Ni He Forest Bureau, China and officers from WWF China Northeastern Center.

References

Alumur, S. A.; Yaman, H.; and Kara, B. Y. 2012. Hierarchical multimodal hub location problem with time-definite deliveries. *Transportation Research Part E: Logistics and Transportation Review* 48(6): 1107–1120.

Breiman, L. 1996. Bagging predictors. *Machine learning* 24(2): 123–140.

Chase, M. J.; Schlossberg, S.; Griffin, C. R.; Bouché, P. J.; Djene, S. W.; Elkan, P. W.; Ferreira, S.; Grossman, F.; Kohi, E. M.; Landen, K.; et al. 2016. Continent-wide survey reveals massive decline in African savannah elephants. *PeerJ* 4: e2354.

Chen, T.; and Guestrin, C. 2016. Xgboost: A scalable tree boosting system. In *SigKDD*, 785–794.

Fang, F.; Nguyen, T.; Ford, B.; Sintov, N.; and Tambe, M. 2015. Introduction to green security games. In *IJCAI*.

Gholami, S.; Ford, B.; Fang, F.; Plumptre, A.; Tambe, M.; Driciru, M.; Wanyama, F.; Rwetsiba, A.; Nsubaga, M.; and Mabonga, J. 2017. Taking it for a test drive: a hybrid spatio-temporal model for wildlife poaching prediction evaluated through a controlled field test. In *ECML PKDD*, 292–304. Springer.

Gholami, S.; Mc Carthy, S.; Dilkina, B.; Plumptre, A.; Tambe, M.; Driciru, M.; Wanyama, F.; Rwetsiba, A.; Nsubaga, M.; Mabonga, J.; et al. 2018. Adversary models account for imperfect crime data: Forecasting and planning against real-world poachers. In *AAMAS*, 823–831. International Foundation for Autonomous Agents and Multiagent Systems.

Gholami, S.; Xu, L.; Carthy, S. M.; Dilkina, B.; Plumptre, A.; Tambe, M.; Singh, R.; Nsubuga, M.; Mabonga, J.; Driciru, M.; et al. 2020. Stay Ahead of Poachers: Illegal Wildlife Poaching Prediction and Patrol Planning Under Uncertainty with Field Test Evaluations. *ICDE* .

Gholami, S.; Yadav, A.; Tran-Thanh, L.; Dilkina, B.; and Tambe, M. 2019. Don't Put All Your Strategies in One Basket: Playing Green Security Games with Imperfect Prior Knowledge. In *AAMAS*, 395–403. International Foundation for Autonomous Agents and Multiagent Systems.

Gurumurthy, S.; Yu, L.; Zhang, C.; Jin, Y.; Li, W.; Zhang, X.; and Fang, F. 2018. Exploiting data and human knowledge for predicting wildlife poaching. In *COMPASS*, 1–8.

Huang, T.; Shen, W.; Gu, T.; Singh, R.; and Fang, F. 2020. Green Security Game with Community Engagement. In *AAMAS*.

Kar, D.; Ford, B.; Gholami, S.; Fang, F.; Plumptre, A.; Tambe, M.; Driciru, M.; Wanyama, F.; Rwetsiba, A.; Nsubaga, M.; and Mabonga, J. 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*, AAMAS '17, 159–167. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Lemieux, A. M. 2014. *Situational prevention of poaching*. Routledge.

McKelvey, R. D.; and Palfrey, T. R. 1995. Quantal response equilibria for normal form games. *Games and economic behavior* 10(1): 6–38.

Nguyen, T. H.; Sinha, A.; Gholami, S.; Plumptre, A.; Joppa, L.; Tambe, M.; Driciru, M.; Wanyama, F.; Rwetsiba, A.; Critchlow, R.; and Beale, C. M. 2016. CAPTURE: A New Predictive Anti-Poaching Tool for Wildlife Protection. In *Proceedings of the 2016 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '16, 767–775. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450342391.

Rasmussen, C. E. 2003. Gaussian processes in machine learning. In *Summer School on Machine Learning*, 63–71. Springer.

Spillane, J. J. 2016. AFRICA'S RHINO POPULATION: PERMANENTLY DECLINING OR SUSTAINABLE? *The Eastern African Journal of Hospitality, Leisure and Tourism.* 3(2): 1–35.

SteadieSeifi, M.; Dellaert, N. P.; Nuijten, W.; Van Woensel, T.; and Raoufi, R. 2014. Multimodal freight transportation planning: A literature review. *European journal of operational research* 233(1): 1–15.

Tambe, M. 2011. *Security and game theory: algorithms, deployed systems, lessons learned*. Cambridge university press.

Wang, Y.; Shi, Z. R.; Yu, L.; Wu, Y.; Singh, R.; Joppa, L.; and Fang, F. 2019. Deep reinforcement learning for green security games with real-time information. In *AAAI*, volume 33, 1401–1408.

Warchol, G. L. 2004. The transnational illegal wildlife trade. *Criminal justice studies* 17(1): 57–73.

Xu, H.; Ford, B.; Fang, F.; Dilkina, B.; Plumptre, A.; Tambe, M.; Driciru, M.; Wanyama, F.; Rwetsiba, A.; Nsubaga, M.; et al. 2017. Optimal patrol planning for green security games with black-box attackers. In *GameSec*, 458–477. Springer.