Predicting Wildlife Trafficking Routes with Differentiable Shortest Paths

Aaron Ferber[1(✉)], Emily Griffin[2], Bistra Dilkina[1], Burcu Keskin[3], and Meredith Gore[4]

[1] University of Southern California, Los Angeles, USA
{aferber,dilkina}@usc.edu
[2] Babson College, Babson Park, USA
egriffin@babson.edu
[3] University of Alabama, Tuscaloosa, USA
bkeskin@cba.ua.edu
[4] University of Maryland, College Park, USA
gorem@umd.edu

Abstract. Wildlife trafficking (WT), the illegal trade of wild fauna, flora, and their parts, directly threatens biodiversity and conservation of trafficked species, while also negatively impacting human health, national security, and economic development. Wildlife traffickers obfuscate their activities in plain sight, leveraging legal, large, and globally linked transportation networks. To complicate matters, defensive interdiction resources are limited, datasets are fragmented and rarely interoperable, and interventions like setting checkpoints place a burden on legal transportation. As a result, interpretable predictions of which routes wildlife traffickers are likely to take can help target defensive efforts and understand what wildlife traffickers may be considering when selecting routes. We propose a data-driven model for predicting trafficking routes on the global commercial flight network, a transportation network for which we have some historical seizure data and a specification of the possible routes that traffickers may take. While seizure data has limitations such as data bias and dependence on the deployed defensive resources, this is a first step towards predicting wildlife trafficking routes on real-world data. Our seizure data documents the planned commercial flight itinerary of trafficked and successfully interdicted wildlife. We aim to provide predictions of highly-trafficked flight paths for known origin-destination pairs with plausible explanations that illuminate how traffickers make decisions based on the presence of criminal actors, markets, and resilience systems. We propose a model that first predicts likelihoods of which commercial flights will be taken out of a given airport given input features, and then subsequently finds the highest-likelihood flight path from origin to destination using a differentiable shortest path solver, allowing us to automatically align our model's loss with the overall goal of correctly predicting the full flight itinerary from a given source to a destination. We evaluate the proposed model's predictions and interpretations both quantitatively and qualitatively, showing that the predicted paths are aligned with observed held-out seizures, and can be interpreted by policy-makers.

A. A. Cire (Ed.): CPAIOR 2023, LNCS 13884, pp. 460–476, 2023.
https://doi.org/10.1007/978-3-031-33271-5_30

1 Introduction

Wildlife Trafficking (WT) broadly impacts biodiversity, human health, economic development, and national security [37]. It encompasses a wide array of species that originate from, and are transported to, supply and demand markets around the world. WT spans over 150 countries and includes more than 37,000 species of fauna and flora [37]. Transnational criminal organizations are known to leverage the increasingly interconnected air transportation network to move illegal wildlife products from source to destination locations, generating $19 billion annually in black market proceeds [18,28,38]. The massive scope, scale, and diversity of wildlife trafficking networks present a complex and dynamic challenge for authorities and researchers trying to understand and interrupt the transiting of illegal wildlife products using detection, interdiction, deterrence, education, or other activities. Stakeholders working to combat wildlife trafficking also face limited social, physical, and financial capital compared to other illicit activities such as drug trafficking. Current practice is to rely heavily on trusted and established personal relationships, "tip-offs" about specific flights, use of specially trained sniffer dogs, and education of airport personnel; these practices can be successful in one-off contexts but lack a desired deterrent effect. Network interdiction models can assist in determining the optimal allocation of scarce resources along known trafficking networks but have yet to be systematically applied to the transiting stage of wildlife trafficking supply chains [17,31]. Data-driven methods for understanding underlying wildlife trafficking patterns could help advance on the ground practice and expand modeling techniques to a novel domain space and are a necessary first step before targeted interdiction allocation can be applied effectively and efficiently.

Recognizing the potential for data-driven methods to dramatically enhance solutions to the problem of wildlife trafficking, multiple sectors have increased their data collection activities. For example, The Convention on International Trade in Endangered Species of Wild Fauna and Flora (CITES) is a global agreement among governments to regulate international trade in species under threat that was established in 1976 and is currently signed by 183 countries and the European Union. TRAFFIC is an organization that was established in 1976 by The World Wide Fund for Nature (WWF) and International Union for Conservation of Nature (IUCN) as a wildlife trade monitoring network to undertake data collection, analysis, and provision of recommendations to inform decision making on wildlife trade. In 2015, the U.S. Agency for International Development (IUCN) established the Reducing Opportunities for Unlawful Transport of Endangered Species (ROUTES) Partnership to bring together transport and logistics companies, government agencies, law enforcement, and conservation organizations to eliminate wildlife trafficking from the air transport supply chain. Importantly, these efforts have contributed to collection and synthesis of a limited but growing global database of illegal wildlife trade seizure data.

Overall, the flight network's widespread use for moving illegal goods, as well as the presence of structured data make it a promising setting for data analysis to help inform defensive measures. Center for Advanced Defense Studies (C4ADS), a nonprofit that is a member of ROUTES, produced in-depth summary analysis of the global wildlife trade flight seizure data from 2009–2017 [38] and 2016–2018 [39] and derived insights based on observed concentration of illegal activity and outliers. Some studies and reports describe traffickers' modus operandi, or factors that may influence their decisions to traffic products through certain ports over others [3,34]. Factors, such as larger airports with higher volume, prevalence of corruption, lower financial costs, and smaller legal penalties, have been shown to possibly be beneficial for traffickers [16]. However, there is limited quantitative research into the factors that impact traffickers' transit choices and their relative importance [22,33,35]. In fact, to our knowledge, predictive models have not been applied to the wildlife trafficking domain. Machine learning models can be instrumental in extrapolating the patterns from the limited seizure data to other airports and routes. They can highlight important factors and their weights to provide insight into traffickers' objectives that can be utilized when making interdiction decisions and predicting trafficker responses.

To this end, in this paper, we formulate wildlife trafficking across the global flight network as a route prediction problem on a graph, synthesize historical seizure data with data that describes airport nodes and flight edges, and propose a maximum likelihood machine learning model that exploits recent developments in differentiable optimization. In particular, we model probabilities of trafficking on each edge in the transportation network as a function of node and edge features, and train the model by comparing the maximum likelihood path (identified by computing the shortest path in log space) to the ground truth paths. We demonstrate the predictive power of our model. We analyze our model's results to understand the discrepancies between our predictions and the ground truth seizure data. By utilizing an interpretable linear model with respect to input features, we are also able to provide feature importance insights.

A key area of concern in combating WT is the convergence of multiple forms of illicit trade [14,35]. Convergence can take a variety of forms. For instance, revenue from WT activities can fund arms trafficking. Additionally, the people, countries, and transit routes used for various forms of trafficking can substantially overlap due to factors that are mutually beneficial. Convergence has long been an area of concern but the amount of scientific, quantitative, evidence for convergence is still limited [15]. Our work makes a step towards quantifying the scale and impact of convergence by directly incorporating measures of other illicit activities at given locations as features when predicting wildlife trafficking paths. Understanding the impact of other illicit activities on the path probabilities of traffickers provides a quantitative measure of geographic convergence.

2 Related Work

The overall problem of learning route choices may be considered an inverse optimization problem, where we are given "solutions" to optimization problems and we want to identify what optimization parameters yields those observed solutions as optimal [1]. Indeed, previous work in trajectory prediction has modeled hidden

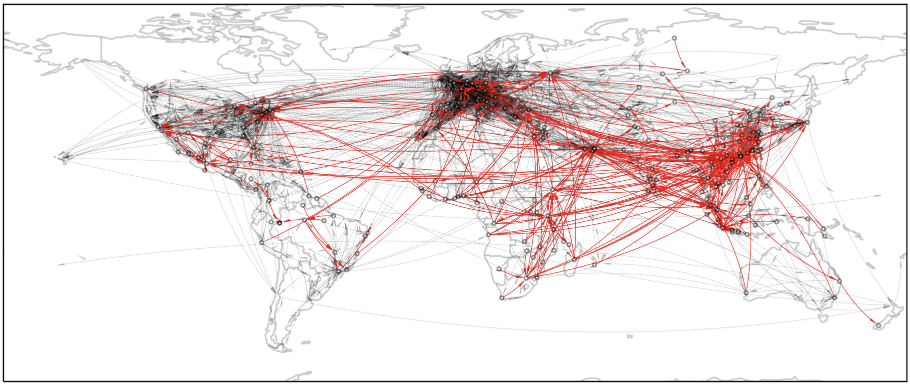

Fig. 1. Visualization of itineraries with historical seizures in red as well as a subset of the global flight network in grey. (Color figure online)

latencies for travel networks by solving an inverse shortest path problem [42], or learning transportation preferences for a road network which results in a given traffic flow on the network [11]. The area of trajectory prediction [10,12,29,43] aims to predict paths for individuals and thus tend to assume access to the start location, or continually updating sequence of locations, and try to predict the rest of the trajectory that the person will take. However, in our case, we have generally-known source and destination pairs and try to understand what are the most likely paths that traffickers will take without continuously updating information.

Recent work in the machine learning literature has investigated how to integrate optimization solvers as differentiable components in machine learning pipelines. This effectively allows the model designer to state that the model predictions will be used downstream by a structured optimization problem which will output an optimal solution to a problem with given predicted inputs. The seminal OptNet paper [2] introduces the quadratic optimization program as a differentiable layer for use in deep learning pipelines, by implicitly differentiating through the KKT optimality conditions, with follow-up work extending the approach to linear programs [40]. In a different vein, researchers investigated differentiating through blackbox optimizers [26] and differentiating through maximum likelihood estimation which can represent the optimal solution to a mathematical program [24]. Our approach directly builds off of [26] and leverages empirical insights in order to speed up gradient computation. Lastly, several approaches for smart predict then optimize have been proposed which compute subgradients of the optimal solution with respect to the inputs in order to train the predictive model [9]. This smart predict then optimize area has work on applicable theoretical guarantees and integration with decision trees [4,8].

Prior work has successfully used machine learning in the context of wildlife poaching in conservation areas, but poaching is only the "first" step in the wildlife trafficking supply chain [13,23,41]. Poaching-oriented approaches consider classification models that predict the likelihood of snare detection at a given spatial location to inform ranger patrolling efforts at the sourcing of wildlife. While these works demonstrated the ability to predict poacher behavior at each pixel of a given conservation area, here we address the global wildlife trade problem of learning trafficker route choices on the broader international air transportation network.

3 Flight Itinerary Prediction Formulation

We formulate the problem of predicting trafficker flight paths connecting a given source airport s and intended destination d airport as a supervised learning problem of predicting a path from s to d on a flight network represented as a directed graph G. The flight network G represents airports as nodes and the flights between them as directed edges. We augment the flight network with WT-related features ϕ on both nodes ϕ^v and edges ϕ^e. We collect N ground-truthed trafficker paths $\mathcal{D}^\pi = \{\pi_{s_i,d_i}\}_{i=1}^N$ from centralized databases of seizure reports. These reports contain the traffickers' intended itineraries between fixed source s_i and destination d_i. We encode these WT itineraries π_i as paths in the flight network, representing them as either a sequence of airport nodes or flight edges as needed.

Our data sources, collection, and synthesis are described in the section "Data Sources". To get a sense of the magnitude of the problem at hand, we visualize the observed trafficker paths as well as 20% of the full flight network in Fig 1. We subsample due to the density of the global flight network consisting of 14,118 flight edges connecting 1,933 airport nodes, rendering the image unreadable otherwise.

Formally, we aim to train a model that correctly predicts the observed structured path π_i given the input source s_i, destination d_i, flight network G, and features ϕ.

3.1 Predictive Model: Edge Transition Estimator

In order to predict full flight paths from features on just edges and nodes, we cannot simply predict how likely any individual path is, as the number of possible simple paths is exponentially large in the size of the network. Instead, we consider predicting a probability for each edge which then can be used to compute path likelihood.

We propose an approach for modeling the path prediction problem by predicting "transition" probabilities, or probabilities on which flight edges a trafficker may take to exit a given "current" node. Since our setting requires a simple model that can be easily handed off to domain experts and deliver actionable insights for interdiction, we forego complex architectures in favor of a simplistic predictive model. This models the trafficker as taking a biased random walk from the source airport to the destination airport on the flight network where our model learns the biased probabilities given edge and node features. With this transition probability modeling approach, we can compute the probability of taking any given source-destination path as being the product of individual edge probabilities.

Formally, we model the problem as finding the probability $P(i, j)$ of using a directed edge (i, j) to leave a starting node i. Here, probabilities on all edges leaving a given node i sum to 1. We use a parametrized model m, with parameters θ, to obtain probability estimates given the relevant features i.e. $\hat{P}(i, j) = m\left(\phi^e_{i,j}, \phi^v_i, \phi^v_j; \theta\right)$. For notational simplicity, we consider the feature vector for a given edge to be the concatenation of edge-specific features, origin features, and destination features $\phi_{i,j} = \left[\phi^e_{i,j}, \phi^v_i, \phi^v_j\right]$. The edge probability prediction model limits the number of trained parameters to prevent overfitting. This parameter sharing means that the same model is used to predict which flights will be taken out of an airport whether it is Addis Ababa or Charles de Gaulle. Furthermore, by predicting edge probabilities from edge and node features, we can understand how these features impact our model's estimates and thus better understand what factors may be driving wildlife trafficking. Hence, in our experiments we use a linear model relating the features to the predicted probabilities to ensure that the resulting model is interpretable.

We denote the set of edges leaving i as $\delta(i)$, and fully specify our linear model as making predictions on each edge as computing logits with a linear model, and using a softmax to normalize the edge logits based on the flight origin node to ensure that the outgoing probabilities sum to one. Mathematically our probability prediction model is described in Eq. 1.

$$\hat{P}(i, j) = m\left(\phi^e_{i,j}, \phi^v_i, \phi^v_j; \theta\right) = \frac{\exp\left(\theta^T \phi_{i,j}\right)}{\sum\limits_{j' \in \delta(i)} \exp\left(\theta^T \phi_{i,j'}\right)} \tag{1}$$

Our formulation ensures that the output probability estimates are a differentiable function of the parameters θ to be trained using standard deep learning libraries like pytorch [25]. Additionally, we experimented using a 3-layer multi-layer perceptron (MLP) as well as gradient-boosted decision trees but found poor generalization of the MLP and the gradient-boosted decision trees performed on par with our linear model so we opted for the linear model as it was interpretable with no drawback in performance.

With the given formulation, the probability of a path $P(\pi)$ is the product of individual edge probabilities $\Pi_{(i,j)\in\pi}\hat{P}((i, j)|i)$. Furthermore, we can identify the model's highest-likelihood path by finding a shortest path with edge weights corresponding to the negative log probability. A path minimizing the sum of negative log probabilities is a path that maximizes the sum of log probabilities which, due to the logarithm's product rule and monotonicity, is a maximum likelihood path. The goal now is to find model parameters θ such that the observed trafficking paths π have the highest likelihood.

At deployment time, this edge transition model will enable us to identify easily the highest-likelihood path by solving a shortest path problem in log probability space, obtain other highly-likely paths by identifying other near-optimal solutions, and allows us to easily evaluate the likelihood of any other alternative path.

3.2 Model Training: Path-Integrated Learning

Given that we want to predict full paths in the flight network, we propose training the parameters θ to directly minimize differences between the predicted highest-probability path and observed trafficking paths. We consider a differentiable pipeline and loss function that directly aligns model training with the problem of recovering the ground truth path, and can be optimized using gradient descent.

Using the above definition of our edge transition probability estimator, we express model training as solving the optimization problem in Eq. 2 which minimizes the expected Hamming loss between a given ground-truth path $\pi_{s,d}$ with corresponding source s and destination d against the highest-likelihood path $\hat{\pi}_{s,d}$ predicted by the model connecting that source to that destination. The highest-likelihood path is computed by Single Source Single Destination shortest path solver (SSSDSolver) over the negative log of predicted transition probabilities $\hat{P}$. Transition probabilities $\hat{P}$ are computed according to Eq. (1).

Ultimately, to train the model we compute gradients for the model parameters via backpropagation of the hamming loss to the predicted highest-probability path $\hat{\pi}$, back to the predicted transition probabilities $\hat{P}$, and then to the model parameters θ.

$$\min_{\theta} E_{\pi_{s,d}} \left[H\left(\pi_{s,d}, \text{SSSDSolver}\left(-log\left(\hat{P}\right); s, d\right)\right)\right] \tag{2}$$

For completeness, we can define the single source single destination shortest path solver in Eq. (3) as finding the path minimizing the sum of weights on edges used in the path π, which in our case are negative log probabilities.

$$\text{SSSDSolver}(w; s, d) = \underset{\pi_{s,d}}{\arg\min} \left(\textstyle\sum_{(i,j)\in\pi_{s,d}} w_{i,j}\right) \tag{3}$$

Here we can use any off-the-shelf shortest path solver without worrying about negative edge weights since the probabilities are all between 0 and 1 (exclusive), so the negative log of the probabilities are all positive values. In practice, we use Dijkstra's shortest path algorithm. Note that the forward pass to get predicted path $\hat{\pi}$ is the same approach we use for determining the highest-likelihood path, thus aligning our model's training with the overall deployment pipeline of correctly identifying the full path.

In order for us to use gradient descent to train our model parameters, we need to ensure that all steps from the model predictions to the loss evaluation are differentiable so that gradients may be easily computed using chain rule. All of the components except for the SSSDSolver are readily differentiable functions available in Pytorch [25], as a result we need to define a backward pass for the shortest path solver to enable model training.

Using the formulation enabling differentiation of blackbox solvers proposed in [26], we make our forward and gradient update explicit below. In the forward pass, we simply solve the shortest path problem and cache the solution $\hat{\pi} := \text{SSSDSolver}(w; s, d)$. The backward pass itself expects incoming gradients from the loss layer, and returns outgoing gradients with respect to the input edge costs w. Overall, the intention of the gradient is to give an indication of what changes in the edge costs w will produce the desired change in the returned path to minimize the loss and better align the path with the ground truth solution. The method for differentiating blackbox solvers introduced in [26] essentially perturbs the input objective coefficients w in the direction of the gradient to find a "locally-improved" solution. It then computes the gradients as the difference between the resulting "locally-improved" solution and the previously predicted solution. When used in conjunction with the hamming loss, the "locally-improved" objective coefficients are simply the input objective coefficients with a given amount increased or decreased depending on whether the decision component, such as the edge usage, should be used or not. In order to specify the degree that the input costs should be perturbed, the authors use a hyperparameter λ which determines the degree to which the weights w should be perturbed in the desired direction. Formally, in the backward pass, we are given input gradients $\nabla_{\hat{\pi}} L$ of the loss with respect to the shortest path $\hat{\pi}$. We compute improved edge weights $w' = w + \lambda \nabla_{\hat{\pi}} L$. Then we re-solve the problem with improved edge weights to find a better solution $\pi' = \text{SSSDSolver}(w'; s, d)$. Finally, we compute gradients of this layer as $-\frac{1}{\lambda}(\hat{\pi} - \pi')$.

In our setting, this method corresponds to solving the shortest path problem with perturbed weights where weight is slightly decreased on edges that should appear in the ground truth solution and slightly increased on edges that aren't in the ground truth solution. The gradient that is passed back to the edge costs is the difference between the predicted path and the "locally-improved" path. Intuitively, the approach aims to decrease cost on edges that should be in the locally-improved short path but aren't in the predicted path, and increase cost on edges that are in the outputted shortest path but don't appear in the locally-improved path. Additionally, in our initial experiments, we found that performant values of λ were large enough so that the weight perturbation eclipsed the initial weights themselves, meaning that overall the "locally-improved" solution was simply the ground truth solution. As such, to cut the number of solves down by half, we simply used the ground truth solution path π as the "locally-improved" solution.

Note that this approach is akin to updating the gradients such that it scores the ground truth solution π to have better objective value than the predicted solution $\hat{\pi}$. Additionally, in this scenario we consider that the path π is encoded as a 0–1 vector with a given entry indicating whether edge (i, j) is used in the path or not. As such, the weight vector is updated by the difference between path solutions.

3.3 Model Training: Edge-Myopic Learning

We compare our path-integrated learning method with an approach that is trained to minimize the Kullback-Leibler (KL) divergence [19] between the edge probabilities computed directly from training data P' to the edge probabilities predicted as a function of features $\hat{P}$. This approach focuses on correctly predicting rates at which different edges are used for trafficking in the ground truth rather than looking at full paths, and is a slight variant of baseline approaches in previous work [5] that is adapted for our setting where we have known source and destination locations, as well as network features. Previous work estimates the transition probabilities between different locations, and here we estimate these transitions with a logistic regression model to obtain a model of how the features are related to the observed transitions. Using raw training data, we estimate transition probabilities $P'(e)$ as the number of times that a given flight e is used for trafficking divided by the number of times that the source airport is used for trafficking. The predictive model's parameters are then trained to closely match these transition probabilities based on the given features. Given probability predictions $\hat{P}$ and data-driven estimates $P'(e)$, the KL divergence is $KL\left(\hat{P}||P'(e)\right) = \sum_e \hat{P}_e \log \frac{\hat{P}_e}{P'(e)}$. Overall, the Edge-Myopic learning trains the parameters θ to minimize this edge-level KL divergence.

4 Data Sources

Centralized and comprehensive data sources are critical for combating wildlife trafficking [16]; however, they are often lacking in practice, complicating the application of models to different domains. In our experiments, we leverage data regarding wildlife trafficking seizures, flight pricing, available flights, and indices of general crime prevalence and resilience infrastructure. The global flight network was collected from OpenFlights.org which hosts open-source information about airports, routes, and flights. The data was last updated on January 2017, and we have manually added several airports and routes to ensure that we can place as many seizure records on the flight network as possible. Overall, this dataset allows us to construct a flight network of 1,933 airport nodes connected by 14,118 flight edges.

For each flight edge, we record the distance and collect flight pricing estimates using the Skyscanner API [30]. Since flight pricing depends on several components such as the amount of time before the flight, we collect prices for all flight routes one month in advance. For each flight edge, we used the API on October 14, 2021 to request flight quotes for November 2021. The API did not return valid responses for several airport pairs due to no valid flight plans existing in the database accessed by the API which we determined manually from searching google flights. Additionally, we note that data was collected during the coronavirus pandemic impacting flight availability, as historical data was not available.

Each airport is associated with its country's metrics reported in the Global Organized Crime Index for 2021 [7], the first year the indices were published by

Table 1. Node and edge features of the flight network. Features in Bold were selected by recursive feature elimination.

NODE FEATURES	
METADATA	
Population Flight Count	CITES membership
GITOC - CRIMINAL MARKETS	
Criminal Markets (Average) **Human Trafficking** Human Smuggling **Arms Trafficking** **Flora Crimes** Non-Renewable Resource Crimes	**Fauna Crimes** **Heroin Trade** Cocaine Trade Cannabis Trade Synthetic Drug Trade
GITOC - RESILIENCE	
Anti-Money Laundering Systems Political Leadership And Governance Govt. Transparency & Accountability Economic Regulatory Capacity Victim & Witness Support Judicial System And Detention	Resilience (Average) Territorial Integrity Law Enforcement International Cooperation National Policies & Laws Prevention
GITOC - CRIMINAL ACTOR	
Criminal Actors (Average) Mafia-Style Groups **Criminal Networks**	State-Embedded Actors Foreign Actors Non-State Actors
EDGE FEATURES	
Price	Distance

the Global Initiative Against Transnational Organized Crime (GITOC). These indices represent expert opinion of a country's relationship with various forms of organized crime, including the prevalence of different criminal actors, strength of resilience resources, and presence of criminal markets. These indices score countries from 1 to 10 based on 5 rounds of anonymous and independent expert reviews in 2020. We also add information about whether the airport's country is a member of CITES, the city's population, and the number of flights that serve the given airport. The node and edge features we collected are summarized in Table 1.

We obtained seizure data from the Wildlife Trade Portal (WTP) [36] through which TRAFFIC provides historical seizure data with detailed records like intended itinerary (source, destination, transit points), trafficked wildlife, trafficker details, and legal outcomes. In total, we accessed 1,067 records between

2017 and 2021 to synthesize a dataset of 454 itineraries of wildlife trafficking. Only 362 of the 1,933 airport nodes in the global flight network are used by traffickers in the historical seizure data, highlighting the data sparsity. Furthermore, in terms of the paths themselves, the data is biased towards shorter paths, having 1-hop, 2-hop, 3-hop, and 4-hop paths making up 60.6%, 24.2%, 15%, and 0.2% of the data respectively.

Seizure data provides a glimpse of how WT networks operate, alert experts to trends in supply and demand for different species, and point to key locations for deterring wildlife crime [20]. However, it is important to understand the biases in seizure data due to being collected by different law enforcement agencies, using several means of detection, against various criminal agents [6,15]. As a result, seizure data not only reflects the criminal network, but also the defensive resources. Nevertheless, seizure data is one of the few tools we have available to peer into WT networks in a scalable manner.

5 Experiments

Table 2. Summary statistics from 10-fold cross-validation of models using either the full set of features or an algorithmically-selected subset. We evaluate two training methods, edge-myopic learning which aims to correctly predict how often individual edges are used, and path-integrated learning which aims to identify the complete intended source-destination path. We report the average performance across folds with 95% confidence intervals.

Training Method	Features	Path recall ↑	Edge recall ↑	Edge precision ↑	Edit distance ↓
Edge-Myopic	Selected	89.6% ± 1.2	83.1% ± 2.6	86.1% ± 1.0	0.115 ± 0.04
Path-Integrated	Selected	92.4% ± 2.7	88.4% ± 4.3	90.5% ± 2.6	0.088 ± 0.03
Edge-Myopic	All	89.2% ± 2.5	82.8% ± 3.5	85.5% ± 3.1	0.113 ± 0.03
Path-Integrated	All	89.1% ± 3.1	82.6% ± 4.0	85.4% ± 3.4	0.113 ± 0.03

5.1 Feature Selection

Feature selection identifies the highest-impact features, limits overfitting, and avoids correlated features. We use recursive feature elimination to iteratively remove the least-useful feature from the current feature set by testing each of them and evaluating the change in 10-fold path recall. Since node features appear twice for a given edge, once for the edge's head and again for the tail, we drop both as needed. The full set and selected features in bold are in Table 1.

5.2 Metrics

We train the models using the Adam optimizer with amsgrad [27], and evaluate the models using 10-fold cross-validation, splitting the dataset by source-destination pair. This produces a prediction for every source-destination path using a model that wasn't trained on information from the given source-destination pair. We evaluate using several metrics at the path and edge level.

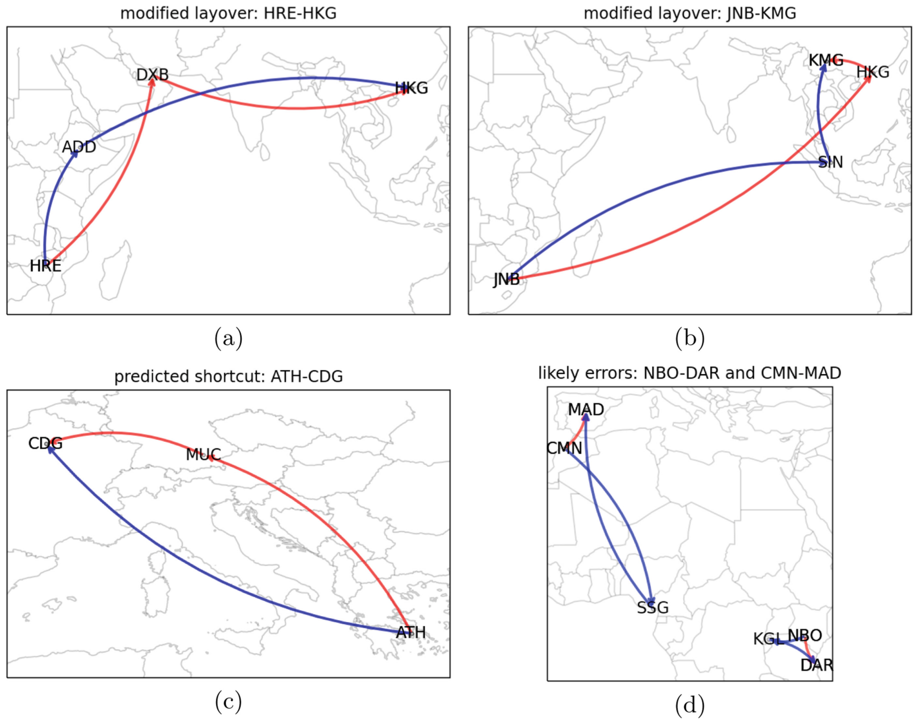

Fig. 2. Visualization of discrepancies between Path-Integrated predicted itineraries in blue and observed itineraries in red. Additionally, domain experts identified two likely errors in Fig. d where our model's predictions are unrealistic. (Color figure online)

Path Recall is the percent of the ground truth paths the model completely predicted correctly. Given the N ground truth paths in the dataset $\mathcal{D}^\pi$, the path recall is $\frac{1}{N}\sum_{\pi_{s,d}\in\mathcal{D}^\pi}\delta\left(\pi_{s,d}=\hat{\pi}_{s,d}\right)$. Here δ is just a 1 if the paths are completely equal (taking the same sequence of edges) and 0 otherwise. Higher values here mean that our model is not likely to miss out on trafficked paths.

Edge Recall is the percent of trafficked edges that our model predicts to have trafficking. Mathematically this is $\left(\sum_{\pi_{s,d}\in\mathcal{D}^\pi}\sum_{e\in\pi_{s,d}}\delta\left(e\in\hat{\pi}_{s,d}\right)\right)/\sum_{\pi_{s,d}\in\mathcal{D}^\pi}|\pi_{s,d}|$. High values here mean that a large proportion of observed trafficked edges are picked up by our model.

Edge Precision measures the percent of edges that our model predicts to have trafficking which did in fact exhibit trafficking in the seizure data. Mathematically this is $\left(\sum_{\pi_{s,d}\in\mathcal{D}^\pi}\sum_{e\in\hat{\pi}_{s,d}}\delta\left(e\in\pi_{s,d}\right)\right)/\sum_{\pi_{s,d}\in\mathcal{D}^\pi}|\hat{\pi}_{s,d}|$. High values here mean that our model's predictions are trustworthy and that domain users can expect that the model's predictions will likely contain trafficking.

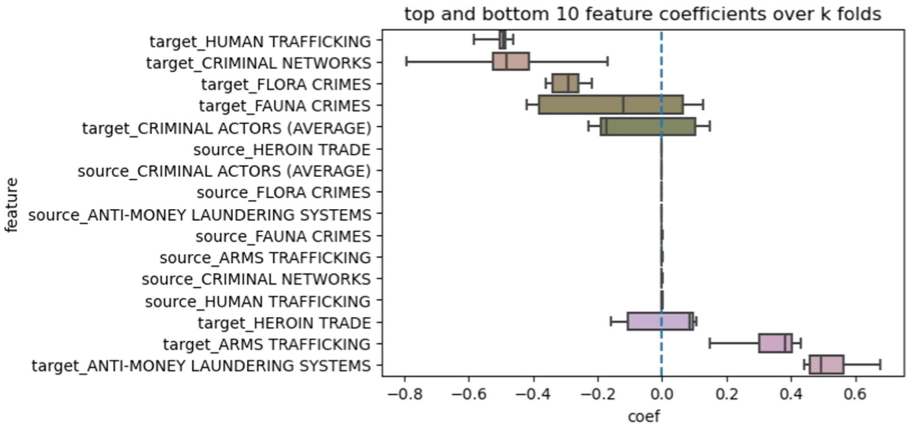

Fig. 3. Feature importance boxplot of the model coefficients across 10 training folds. Positive values suggest higher trafficking rates when the indicator is prevalent and negative values indicate lower rates when the indicator is prevalent.

Edit Distance or Levenshtein distance [21], is the smallest number of "edits" (additions, removals, or substitutions) needed to go from the predicted to ground truth path, considering the itinerary to be a sequence of airports visited. Low edit distance means that the predicted paths are similar to the observed paths.

5.3 Results Discussion

We present numerical results in Table 2, computing the average and standard deviation in performance with 10-fold cross-validation. Given the data size of only 454 itineraries, the differences in performance come from only a few predicted paths. Additionally, both models benefit from feature selection, with feature-selected path-integrated learning improving over edge-myopic learning. The performance of path-integrated learning with feature selection is high in that the models are able to recall 92.4% of the paths completely, 88.4% of the edges, and the predicted edges contain trafficking at a rate of 90.5%. Additionally, breaking the results down by path length, we find that on the 1-hop, 2-hop, 3-hop, and 4-hop paths, path-integrated learning with selected features gets path recalls of 98.9%, 83.1%, 80.7%, and 100%, respectively, and average Levenstein distances of 0.031, 0.169, 0.192, and 0 respectively. On the other hand, across 1-hop, 2-hop, 3-hop, and 4-hop paths, the edge-myopic learning with selected features has path recall of 100%, 83.1%, 58.6%, and 0%, respectively, with Levenstein distances of 0.0, 0.261, 0.314, 2.0, respectively. Across all folds, our model differs with 31 ground truth paths between 25 origin-destination pairs. We note that our performance improvements aren't statistically significant for most metrics, except for edge precision, due to our small sample size. However, given that path-integrated learning gives fewer errors in our low data regime and are promising for future work when more data is available for evaluation. Given the data bias, the predicted alternate routes may contain wildlife trafficking even though it is not present in the ground-truth data. We visualize representative discrepancies between predicted and observed paths in Fig. 2 with predicted paths in blue and observed paths in red. We categorize the discrepancies into 10 origin-destination pairs where our predictions shortcut the observed itinerary by removing stops (Fig. 2a), and 13 cases where our model predicts different layovers than the observed (Fig. 2b, 2c), identified as highly plausible in informal consultations with experts. The two cases where our model predicted additional layovers are visualized in Fig. 2d and are likely errors.

We visualize the path-integrated learning model's feature importance in Fig. 3. Here, positive values mean that high feature values induce high estimated probability, whereas negative values mean that high feature values induce low estimated probability. Overall, the model considers that traffickers are likely to travel to locations with high arms trafficking as well as resilience against money laundering. The convergence between wildlife trafficking and arms trafficking has been documented and has broad implications for interdiction [32]. Additionally, the model predicts that traffickers are less likely to enter regions with high flora crime, criminal networks, or human trafficking. The negative value for the flight destination's flora crimes is interesting and warrants further investigation, and may reflect seizure data bias, or traffickers wanting to flee suspicion. Some features have 0 weight from selecting features based on edge-myopic learning, as well as having correlated features. Ultimately, we propose a model and a training approach, presenting promising results with the best data available so far. More in-depth and robust conclusions about wildlife trafficking route prediction can be made in the future as more complete seizure data and richer feature sets become available, which can leverage our modeling work.

5.4 Conclusion

We approach the problem of predicting wildlife trafficking on the flight transportation network with differentiable optimization. To align our network training with the goal of correctly identifying full paths, we train with a differentiable highest-probability path solver We show that a path-integrated learning model learns over the available airport and flight features with limited training data, slightly improving over edge-myopic learning, and can likely further improve as more seizure data is collected. Lastly, we identify several features that may contribute to traffickers being more likely to take a given path. We hope that our method will help inform interdiction efforts and the study of wildlife trafficking networks, and we intend to use our predictions in conjunction with combinatorial interdiction in future work.

Acknowledgments. All authors were supported by U.S. NSF awards CMMI-1935451; Gore was also supported by IIS-2039951. The information contained herein does not represent the opinions of the U.S. Government or any author affiliations.

References

1. Ahuja, R.K., Orlin, J.B.: Inverse optimization. Oper. Res. **49**, 771–783 (2001). https://doi.org/10.1287/opre.49.5.771.10607
2. Amos, B., Kolter, J.Z.: OptNet: differentiable optimization as a layer in neural networks. In: Proceedings of the 34th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 70, pp. 136–145. PMLR (2017)
3. Arroyave, F.J., Petersen, A.M., Jenkins, J., Hurtado, R.: Multiplex networks reveal geographic constraints on illicit wildlife trafficking. Appl. Netw. Sci. **5**(1), 1–20 (2020). https://doi.org/10.1007/s41109-020-00262-6
4. Balghiti, O.E., Elmachtoub, A.N., Grigas, P., Tewari, A.: Generalization bounds in the predict-then-optimize framework. In: Advances in Neural Information Processing Systems, vol. 32 (2019)
5. Choi, S., Yeo, H., Kim, J.: Network-wide vehicle trajectory prediction in urban traffic networks using deep learning. Transp. Res. Rec. **2672**(45), 173–184 (2018)
6. CITES: https://cites.org/eng ()
7. Global Initiative Against Transnational Organized Crime. The global organized crime index (2021). https://globalinitiative.net/analysis/ocindex-2021/
8. Elmachtoub, A., Liang, J.C.N., McNellis, R.: Decision trees for decision-making under the predict-then-optimize framework. In: International Conference on Machine Learning, pp. 2858–2867. PMLR (2020). https://github.com/rtm2130/SPOTree
9. Elmachtoub, A.N., Grigas, P.: Smart "predict, then optimize". Manage. Sci. **68**(1), 9–26 (2021). https://doi.org/10.1287/mnsc.2020.3922
10. Feng, J., et al.: DeepMove: predicting human mobility with attentional recurrent networks. In: Proceedings of the 2018 World Wide Web Conference, pp. 1459–1468 (2018)
11. Fosgerau, M., Paulsen, M., Rasmussen, T.K.: A perturbed utility route choice model. Transp. Res. Part C Emerg. Technol. **136**, 103514 (2022). https://doi.org/10.1016/j.trc.2021.103514, https://www.sciencedirect.com/science/article/pii/S0968090X21004976
12. Gambs, S., Killijian, M.O., del Prado Cortez, M.N.: Next place prediction using mobility Markov chains. In: Proceedings of the First Workshop on Measurement, Privacy, and Mobility, pp. 1–6 (2012)
13. Gholami, S., et al.: Adversary models account for imperfect crime data: forecasting and planning against real-world poachers (corrected version). In: 17th International Conference on Autonomous Agents and Multiagent Systems (2018)
14. Gore, M.L., et al.: Transnational environmental crime threatens sustainable development. Nat. Sustain. **2**(9), 784–786 (2019)
15. Gore, M.L., Mwinyihali, R., Mayet, L., Baku-Bumb, G.D.M., Plowman, C., Wieland, M.: Typologies of urban wildlife traffickers and sellers. Glob. Ecol. Conserv. **27**, e01557 (2021)
16. Gore, M.L., et al.: Voluntary consensus based geospatial data standards for the global illegal trade in wild fauna and flora. Sci. Data **9**(1), 1–8 (2022)
17. Haas, T.C., Ferreira, S.M.: Finding politically feasible conservation policies: the case of wildlife trafficking. Ecol. Appl. **28**(2), 473–494 (2018)
18. IATA: Combating wildlife trafficking. https://www.iata.org/en/programs/environment/wildlife-trafficking/. Accessed 15 Aug 2022
19. Kullback, S.: Information theory and statistics. Courier Corporation (1997)

20. Kurland, J., Pires, S.F.: Assessing us wildlife trafficking patterns: how criminology and conservation science can guide strategies to reduce the illegal wildlife trade. Deviant Behav. **38**(4), 375–391 (2017)
21. Levenshtein, V.I., et al.: Binary codes capable of correcting deletions, insertions, and reversals. In: Soviet Physics Doklady, vol. 10, pp. 707–710. Soviet Union (1966)
22. Magliocca, N., et al.: Comparative analysis of illicit supply network structure and operations: cocaine, wildlife, and sand. J. Illicit Econ. Dev. **3**(1), 50–73 (2021)
23. Nguyen, T.H., et al.: Capture: a new predictive anti-poaching tool for wildlife protection. In: Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, pp. 767–775 (2016)
24. Niepert, M., Minervini, P., Franceschi, L.: Implicit MLE: backpropagating through discrete exponential family distributions. In: Advances in Neural Information Processing Systems, vol. 34 (2021). https://github.com/nec-research/tf-imle
25. Paszke, A., et al.: PyTorch: an imperative style, high-performance deep learning library. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) Advances in Neural Information Processing Systems, vol. 32, pp. 8024–8035. Curran Associates, Inc. (2019). http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf
26. Pogančić, M.V., Paulus, A., Musil, V., Martius, G., Rolinek, M.: Differentiation of blackbox combinatorial solvers. In: International Conference on Learning Representations (2020). https://openreview.net/forum?id=BkevoJSYPB
27. Reddi, S.J., Kale, S., Kumar, S.: On the convergence of Adam and Beyond. In: International Conference on Learning Representations (2018). https://openreview.net/forum?id=ryQu7f-RZ
28. ROUTES: How the aviation industry transformed to combat wildlife trafficking (2022). https://www.internationalairportreview.com/article/173456/how-the-aviation-industry-transformed-to-combat-wildlife-trafficking/. Accessed 15 Aug 2022
29. Rudenko, A., Palmieri, L., Herman, M., Kitani, K.M., Gavrila, D.M., Arras, K.O.: Human motion trajectory prediction: a survey. Int. J. Robot. Res. **39**(8), 895–935 (2020)
30. skyscanner (2020). https://skyscanner.github.io/slate/
31. Smith, J.C., Song, Y.: A survey of network interdiction models and algorithms. Eur. J. Oper. Res. **283**(3), 797–811 (2020)
32. Spevack, B.: Shared skies convergence of wildlife trafficking with other illicit activities in the aviation industry. In: C4ADS (2021). www.routespartnership.org
33. Stringham, O.C.: Text classification to streamline online wildlife trade analyses. PLoS ONE **16**(7), e0254007 (2021)
34. Stringham, O.C., et al.: Dataset of seized wildlife and their intended uses. Data Brief **39**, 107531 (2021). https://doi.org/10.1016/j.dib.2021.107531. https://www.sciencedirect.com/science/article/pii/S2352340921008076
35. Stringham, O.C., et al.: A guide to using the internet to monitor and quantify the wildlife trade. Conserv. Biol. **35**(4), 1130–1139 (2021)
36. TRAFFIC: Wildlife trade portal (2021). www.wildlifetradeportal.org
37. UNODC: Enhancing the Detection, Investigation, and Disruption of Illicit Financial Flows from Wildlife Crime (2017). http://www.unodc.org/unodc/en/data-and-analysis/wildlife.html
38. Utermohlen, M., Baine, P.: In plane sight: wildlife trafficking in the air transport sector (2018). https://www.traffic.org/publications/reports/in-plane-sight/. Accessed 15 Aug 2022

39. Utermohlen, M.: Runway to extinction: wildlife trafficking in the air transport sector (2020). https://routespartnership.org/industry-resources/publications/runway-to-extinction-report. Accessed 15 Aug 2022
40. Wilder, B., Dilkina, B., Tambe, M.: Melding the data-decisions pipeline: decision-focused learning for combinatorial optimization. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 33, 1658–1665 (2019). https://doi.org/10.1609/aaai.v33i01.33011658
41. Xu, L., Bondi, E., Fang, F., Perrault, A., Wang, K., Tambe, M.: Dual-mandate patrols: multi-armed bandits for green security. In: Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI-21) (2021)
42. Zhang, J., Paschalidis, I.C.: Data-driven estimation of travel latency cost functions via inverse optimization in multi-class transportation networks. In: 2017 IEEE 56th Annual Conference on Decision and Control, CDC 2017, pp. 6295–6300 (2018). https://doi.org/10.1109/CDC.2017.8264608
43. Ziebart, B.D., et al.: Planning-based prediction for pedestrians. In: 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 3931–3936. IEEE (2009)