Learning Opportunistic Adversarial Model on Global Wildlife Trade

Kai Wang
Harvard University
Cambridge, MA
kaiwang@g.harvard.edu

Jeffrey Brantingham
UCLA
Los Angeles, CA
branting@ucla.edu

Milind Tambe
Harvard University
Cambridge, MA
milind_tambe@harvard.edu

ABSTRACT

Global illegal wildlife trade threatens biodiversity and acts as a potential crisis of invasive species and disease spread. Despite a wide range of national and international policies and regulations designed to stop illegal wildlife trade, high profit margins and increasing demand drive a vigorous global illicit trade network. In this paper, we aim to build an adversarial model to predict the future wildlife trade based on the historical trade data. We hypothesize that the majority of *illegal* wildlife trade is opportunistic crime, which is highly correlated to *legal* wildlife trade. We can therefore leverage the abundant legal wildlife trade data to forecast the future wildlife trade, where a fixed fraction of trade volume will reflect the opportunistic wildlife trade volume. To learn a legal wildlife trade model, we propose to use graph neural networks and meta-learning to handle the network and species dependencies, respectively. Lastly, we suggest to incorporate agent-based models on top of our model to study the evolution from opportunistic to more organized illegal wildlife trade behavior.

KEYWORDS

Wildlife trade, opportunistic adversarial model, graph neural networks, meta-learning

ACM Reference Format:
Kai Wang, Jeffrey Brantingham, and Milind Tambe. 2021. Learning Opportunistic Adversarial Model on Global Wildlife Trade. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), London, UK, May 3–7, 2021*, IFAAMAS, 5 pages.

1 INTRODUCTION

Wildlife trade presents a major challenge to environmental sustainability. The high economic value in global wildlife trade drives illegal poaching activity, threatening endangered species and biodiversity. Recent estimates place the value of global illegal wildlife trade at $7-23 billion annually, ranking it third behind only illegal narcotics and weapons trafficking [15, 16, 28]. Various organizations have put considerable effort into regulating and investigating the global wildlife trade. For example, CITES [5], the Convention on International Trade in Endangered Species of Wild Fauna and Flora led by United Nations Environment Programme World Conservation Monitoring Centre (UNEP-WCMC), has collected more than 13 million records of globally regulated legal trade in wildlife of over 34,000 scientific species. CITES is dedicated to understand the surge of global wildlife trade and promotes interventions to protect endangered species. On the other hand, TRAFFIC, a leading non-governmental organization working globally on trade in wild animals and plants, aims ensure that wildlife trade is not a threat to conservation goals. TRAFFIC has conducted many surveys related to illegal wildlife trad [21].

Recently, artificial intelligence has achieved great success in supporting anti-poaching efforts in national parks [6, 31, 32]. There are various issues involved in extending the same techniques to prevent illegal wildlife trade, however. In particular, the network structure underlying trading behavior and the sparsity of trade data present challenges for most of the existing machine learning techniques. In this paper, we focus our study on legal wildlife trade and opportunistic illegal wildlife trade [19, 30], where the smugglers use a volume-based smuggling approach to ship illegal items along with similar legal items, whose outcome is very similar to legal wildlife trade. Specifically, a large fraction of illegal wildlife trade is composed of opportunistic trade [23]. With these, we think forecasting future legal wildlife trade would help us understand future opportunistic illegal wildlife trade.

We explore several interrelated problems. First, we propose to learn an opportunistic adversarial model by predicting the future legal wildlife trade volume, where a fixed fraction of the trade volume becomes opportunistic trade. In this task, we use graph neural networks to handle the underlying network dependencies, where the wildlife trade behavior is highly dependent on the node and edge characteristics. In particular, the past trading data of each particular route are also very informative. Therefore, a model that can properly handle network dependency and time series data should play an important role in predicting future wildlife trade. Second, we note that trade routes for different species are different but correlated, which may relate to differences in the characteristics of the species involved [15, 18], source-side biogeography and sink-side demand. This is an ideal problem for the application of meta-learning where, given a species description, a meta model can adjust and specialize to a particular adversarial model. Meta-learning can often aggregate data from different tasks and thus achieve better performance. Finally, we expect our opportunistic model can be used to help understand the evolution of more organized illegal wildlife trade. Driven by the increasing wildlife trade profit, opportunistic trade can gradually evolve to organized illegal wildlife trade [30]. We hope to use our model to incorporate additional agent-based assumptions to study this wildlife trade evolution [4].

2 WILDLIFE TRADE AND RELATED WORKS

In this paper, we focus on international wildlife trade. Depending on governmental border control policy, wildlife trade can be split into two categories: legal wildlife trade and illegal wildlife trade. Legal

Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, London, UK.

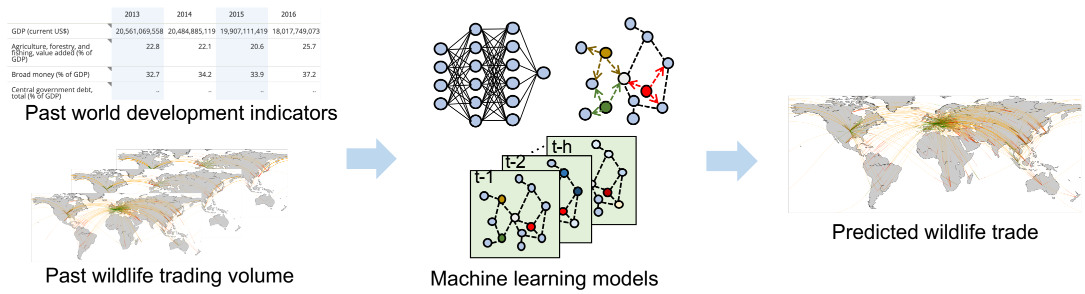

Figure 1: The historical data of global wildlife trade, country development indicators, and network structure are the available features that can be used to predict the future wildlife trade. Different machine learning models, including multilayer perceptron, graph neural networks, spatio-temporal graph neural networks, can be used to fit the historical wildlife trade data. Once the model is trained, we can make prediction on the unseen future wildlife trade.

wildlife trade is under regulatory administration, with the goal of mitigating impacts to biodiversity and the environment. Legal trade is generally well documented in data. Illegal wildlife trade, by contrast, is not under administration, which may lead to greater environmental impact. Illegal trade is also poorly documented in data due to the efforts of smugglers to conceal their activity.

2.1 Agent-based Adversarial Models on Illegal Wildlife Trade

International illegal wildlife trade involves individuals, groups and/or organizations transporting animals or plants, both living and their derivative parts, across international borders in contravention of national and/or international laws. Here we discuss adversarial models of smugglers that differ by the level of rationality of the actors involved.

2.1.1 Fully Rational Adversary. In the most extreme case, smugglers can conduct surveillance to understand the risk of each route, and optimally select the one route with the lowest risk and highest profit. The problem can be formulated as a network security game [24], an extension of the general Stackelberg security game [22]. In network security games, the defender selects a mixed strategy to allocate checkpoints on a network first, while the adversary chooses an optimal path after observing the defender's mixed strategy. Many algorithms [10, 11] were proposed to find the optimal defense strategy in network security games. However, the fully rational assumption is often too pessimistic. Smugglers are not omniscient and often follow learned routines that persist in spite of defender strategies.

2.1.2 Boundedly Rational Adversary. Instead of assuming a perfect adversary, adversaries can have their own unknown preferences, which can be learned from observed historical interactions between the defenders and the adversaries. Various boundedly rational models [7, 27, 33] were proposed to model the bounded rationality. However, the major challenge of fitting a boundedly rational model in the present case is the lack of sufficient illegal wildlife trade data. Not only are past defense policies usually confidential, but data on past illegal smuggling trajectories are hard to collect.

2.2 Opportunistic Illegal Wildlife Trade

Interestingly, various studies suggest that most wildlife smugglers adopt strategies close to normative behavioral routines [3, 4, 15, 18, 19, 28, 30]. This observation motivates us to study opportunistic illegal wildlife trade as a proxy for hard-to-observe illegal wildlife trade. Specifically, we propose to use historical legal wildlife trade data over time to build an adversarial model to predict the future legal wildlife trade, which also serves as an estimate of opportunistic illegal wildlife trade over time. We expect to be able to transfer the knowledge gained from the study of legal, opportunistic wildlife trade to more sophisticated organized illegal wildlife trade.

3 PROBLEM STATEMENT

The main goal is to predict the future wildlife trade quantity of each species at every available cross-country route using the information of historical trading quantities, network structure, and various country indicators. To accommodate the network structure, we use $G = (V, E)$ to denote the entire physical network worldwide, where each node $v \in V$ represents a country, and a directed edge $e = (u, v) \in E$ connecting two countries refers to the trading direction from country u to country v. We only use the edges that have been used as a trading route in the dataset as our edge set E.

In particular, to accommodate the difference between countries and time periods, we assume there is a fixed-length feature $\xi_{v,t} \in \mathbb{R}^k$ associated with each country $v \in V$ at each time stamp $t \in T$. These temporal country-dependent features can be used to identify countries with similar social and economic conditions, which very likely would affect the demand volume and trading decisions. We use $y_{s,e,t} \in \mathbb{R}_{\geq 0}$ to denote the trading volume of species $s \in S$ at route $e \in E$ at time t. Our goal is to use all the information prior to time t, but not the information at time t, to predict the trading volume at time t. The flowchart is shown in Figure 1. For each species $s \in S$, we aim to find the best model Φ^s_w parameterized by w by minimizing the expected loss over the past time $t \in T$ and available routes $e \in E$:

$$E_{t \in T, e \in E} \left[\text{loss}(\Phi^s_w(e, y_{<t}, \xi_{<t}), y_{s,e,t}) \right] \quad (1)$$

We assume our predictive model Φ^s_w takes the historical country features $\xi_{<t}$ and past trading quantities $y_{<t}$ as inputs. For each

feasible edge $e \in E$, the model outputs a value as the predicted future trading volume. At test time, we can use the trained model to predict the trade quantity at future unseen time horizon.

Given the disparity between the trading behavior of different species [15], we learn and maintain a predictive model Φ_w^s for each species $s \in S$. In Section 6.2, we will discuss the potential of applying meta-learning to learn a single meta model to fit all the species.

4 DATASETS

Here we describe two main datasets that we use as our country features and wildlife trade quantities.

4.1 World Development Indicators

The World Development Indicators [1] is a dataset maintained by The World Bank. It is a compilation of relevant, high-quality, and internationally comparable statistics about global development. It includes different categories of statistics of each country, including more than 1400 indicators related to poverty, equality, people, environment, economics, markets, and global links. We use the annual statistics of all these country indicators as our country features.

4.2 Convention on International Trade in Endangered Species (CITES)

The CITES dataset [5] contains more than 13 million records of trade in wildlife of over 34,000 scientific species since 1975. CITES is commonly used to study and monitor the level of international trade [2, 20]. Each data entry contains the taxon of the particular trading species, time period (year), export country, import country, and quantity with unit. We can use these information to construct a trading route map for each species $s \in S$ in a given year $t \in T$, where each edge $e \in E$ is associated with a trading volume $y_{s,t,e}$. We adopt the data collected from 1980 to 2018. All the past wildlife trade volume are used as our edge features, while the trade volume in the next year is used as the target value we want to predict.

5 PREDICTIVE MODELS

Since our input data is graph-dependent and a time series data, we provide the description of various predictive models below.

5.1 Multilayer Perceptron

The most naive model we can use here is multilayer perceptron. We can concatenate the feature of the importing and exporting countries in the last time period, and the past trade volume for the route as our route feature. Therefore, for each route $e = (u, v) \in E$ and each time t, the fixed-length features are composed of $\xi_{u,t-1}$, $\xi_{v,t-1}$, and $\{y_{e,t-i}\}_{i \in \{1,2,\cdots,h\}}$. We can fit a multilayer perceptron model to predict the current trade volume $y_{e,t}$.

5.2 Graph Neural Network

To leverage the network structure of the wildlife trade problem, we can use graph neural networks (GNNs) [29, 36] as our predictive model to resolve the graph dependency. The message passing process implemented in GNNs can propagate the features of nodes and edges to their neighbors in multiple hubs away with a nonlinear activation function [9, 13] or an attention mechanism [25] to aggregate information. For example, we can use the country development indicators in the last time period as our node feature ξ_u, and the past 5 years trade volume and some route properties, e.g., distance, as our edge feature η_e. We can feed the entire network with edge set E being the available routes and the features into a GNN. The GNN performs graph convolutions to propagate node and edge features to their neighborhood and eventually outputs vectors $\hat{\xi}_u, \hat{\eta}_e$ for each node $u \in V$ and edge $e \in E$, serving as a compact embedding of the knowledge of neighborhood structure. We can concatenate the compact embeddings $\hat{\xi}_u, \hat{\xi}_v$ of node u, v and the corresponding edge embedding $\hat{\eta}_e$] and feed them into a multilayer perceptron to predict the corresponding trade volume.

5.3 Spatio-temporal Graph Neural Network

To incorporate the temporal information of country development indicators and wildlife trade, it is better to include the entire time series data of the last few time periods as the input features of our model. There is a similar challenge in traffic prediction [14, 35], where both the network structure and the time series are crucial for making prediction at the next time frame. In particular, Yu et al. [34] proposed Spatio-Temporal Graph Convolutional Networks (STGCN) to include temporal convolution layers in graph neural networks to combine temporal and spatial information. STGCN achieves great success in traffic prediction against the existing baselines. We expect that incorporating both the temporal and spatio features can achieve better performance.

6 FUTURE DIRECTIONS AND CHALLENGES

In this section, we summarize some future directions and the associated challenges of forecasting future wildlife trade.

6.1 Graph Neural Networks Structure

To apply graph neural networks on wildlife trade prediction, the first challenge is how to properly handle node features, edge features, and time series data. In the literature, various edge-dependent graph neural networks [8, 12] were proposed to handle the challenge of edge features and node features. On the other hand, Yu et al. [34] proposed spatio-temporal graph neural networks to handle the time series node features, where they added additional temporal convolutional layers to GNN in order to process and aggregate the time series data. However, the absence of edge features makes it difficult to fit into our problem, where our edge features play an important role in making prediction because the past trade volume is highly correlated to the future trade volume. Therefore, incorporating both the edge features and time series data properly is a crucial challenge in predicting future wildlife trade.

6.2 Meta-learning Across Different Species

Our second challenge is the dependency on trade in different species Moreto and Lemieux. In the previous sections, we suggested to build a model for each species to specialize the trading pattern. However, this approach lets go of the opportunity to leverage correlations between species. A better solution is to build up a single meta model, with the species description used as an input to adjust predictions. This direction is known as meta-learning [17, 26]. In our wildlife trade prediction problem, we have a limited amount of data per

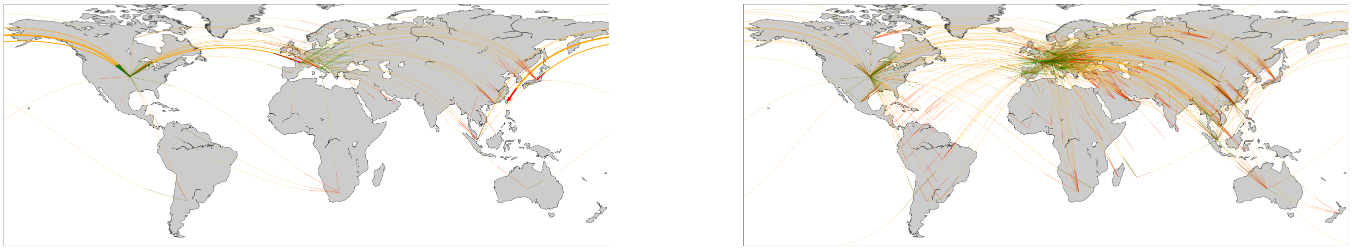

(a) Alligator Mississippiensis trade distribution in 1990

(b) Alligator Mississippiensis trade distribution in 2018

Figure 2: Visualization of the trading route change of American alligator (Alligator Mississippiensis) from 1990 to 2018 in CITES dataset. The routes are plotted with line width proportional to the routes' relative frequency, where green and red arrows refer to the exporting and the importing ends, respectively. The major habitat of American alligator is located in the southeastern United States. We can see many major routes going from United States to some European and Asian countries for redistribution purpose in 1990. In 2018, although the major routes remain the same, we see a more diverse route distribution and a more widespread demand.

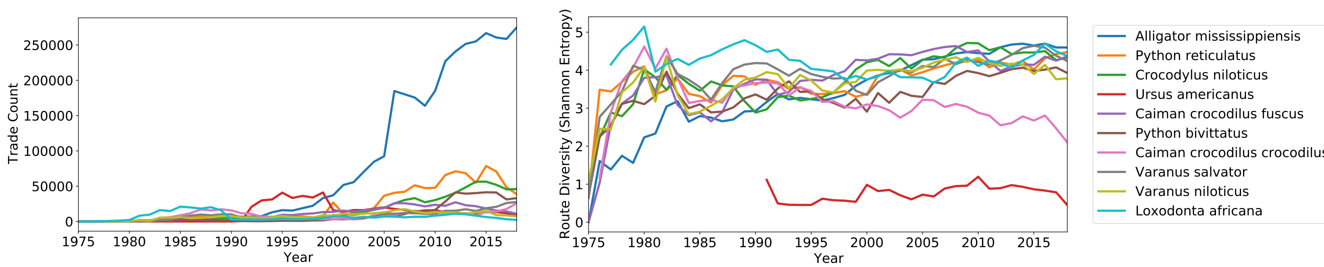

(a) Annual trade counts of various wildlife species from 1975 to 2018.

(b) Visualization of the annual route diversity using Shannon entropy. The larger the entropy is, the more diverse the route usage is.

Figure 3: Visualization of the trade counts and route diversity of the top-10 popular trading species in CITES. We can see that there is a significant growth in terms of the trade counts across all wildlife species. The diversity of route choices benefits from the development of air transportation. However, the diverse shipping route choices can also make administration harder.

species. With the use of meta-learning, we can combine the data from all species and learn the model collectively, which can resolve the issue of data deficiency.

6.3 Non-opportunistic Adversarial Models and Intervention Effect

Although the majority of illegal wildlife trade is potentially opportunistic, when a specific intervention is imposed, trade behavior could change drastically. We still need to understand how the adversary responds to different potential interventions. A first step is to check whether our model learned from the CITES database different trade behaviors between less and more endangered species. This can provide insight on trading behavior shift under different levels of regulation. A more sophisticated non-opportunistic adversarial model, with the opportunistic model as the prior, and intervention response model can be built to accommodate the non-opportunistic behavior and the effect of intervention.

7 VISUALIZATION AND DISCUSSION

In Figure 2, we visualize the networks underlying legal trade in 1990 and 2018 of American alligator. The main trading source of American alligator is the United States. We can see a few clear hubs in Europe in 1990 to redistribute alligator to other countries, while the routes became sparser in 2018 due to the expansion of air transportation and more demand for alligator across different countries. A similar route distribution shift can also be found across all species, but they are highly dependent on the species habitats and local culture. This suggests that there is a potential for machine learning to help predict the complex future trade distribution.

In Figure 3(a) and 3(b), we visualize the trade volume and route diversity of the top-10 traded species. Trade volume and route diversity increased for most species. This suggests a globally increase in wildlife trade issue and an evolving sophisticated trade network. Machine learning models can help by properly handling the complex historical information and network structure.

8 CONCLUSION

Forecasting future wildlife trade is an important interdisciplinary direction for conservationists and computer scientists. We summarize some challenges with proposed machine learning solutions. We hope building an opportunistic adversarial model can help predict and control the prevalent global wildlife trade, and hopefully foster more agent-based adversarial modeling studies in wildlife trade.

REFERENCES

[1] World Bank. 2010. *World development indicators 2010.* The World Bank.
[2] Michal Berec, Lucie Vršecká, and Irena Šetlíková. 2018. What is the reality of wildlife trade volume? CITES Trade Database limitations. *Biological Conservation* 224 (2018), 111–116.
[3] Martin Bouchard and Carlo Morselli. 2014. Opportunistic structures of organized crime. *The Oxford handbook of organized crime* 1 (2014), 288–302.
[4] Steven Broad, Teresa Mulliken, and Dilys Roe. 2003. *The nature and extent of legal and illegal trade in wildlife.* Vol. 3–22. Earthscan Publications, London, 3–22.
[5] UNEP-WCMC CITES. 2016. CITES trade statistics derived from the CITES Trade Database. *UNEP World Conservation Monitoring Centre, Cambridge, UK ([WWW Document]. Conv. Int. Trade Endanger. Species wild Fauna Flora. URL http://trade. cites. org/(accessed 4.16. 16))* (2016).
[6] Fei Fang, Peter Stone, and Milind Tambe. 2015. When Security Games Go Green: Designing Defender Strategies to Prevent Poaching and Illegal Fishing.. In *IJCAI.* 2589–2595.
[7] Benjamin Ford, Thanh Nguyen, Milind Tambe, Nicole Sintov, and Francesco Delle Fave. 2015. Beware the soothsayer: From attack prediction accuracy to predictive reliability in security games. In *International Conference on Decision and Game Theory for Security.* Springer, 35–56.
[8] Liyu Gong and Qiang Cheng. 2019. Exploiting edge features for graph neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition.* 9211–9219.
[9] William L Hamilton, Rex Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. *arXiv preprint arXiv:1706.02216* (2017).
[10] Manish Jain, Vincent Conitzer, and Milind Tambe. 2013. Security scheduling for real-world networks. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems.* 215–222.
[11] Manish Jain, Dmytro Korzhyk, Ondřej Vaněk, Vincent Conitzer, Michal Pěchouček, and Milind Tambe. 2011. A double oracle algorithm for zero-sum security games on graphs. In *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 1.* 327–334.
[12] Jongmin Kim, Taesup Kim, Sungwoong Kim, and Chang D Yoo. 2019. Edge-labeling graph neural network for few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition.* 11–20.
[13] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
[14] Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. 2017. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. *arXiv preprint arXiv:1707.01926* (2017).
[15] William D. Moreto and A. M. Lemieux. 2015. From CRAVED to CAPTURED: Introducing a Product-Based Framework to Examine Illegal Wildlife Markets. *European Journal on Criminal Policy and Research* 21, 3 (2015), 303–320. https://doi.org/10.1007/s10610-014-9268-0
[16] Christian Nellemann, Rune Henriksen, Arnold Kreilhuber, Davyth Stewart, Maria Kotsovou, Patricia Raxter, Elizabeth Mrema, Sam Barrat, et al. 2016. *The rise of environmental crime: a growing threat to natural resources, peace, development and security.* United Nations Environment Programme (UNEP).
[17] Alex Nichol, Joshua Achiam, and John Schulman. 2018. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999* (2018).
[18] Stephen Pires and Ronald V. Clarke. 2012. Are Parrots CRAVED? An Analysis of Parrot Poaching in Mexico. *Journal of Research in Crime and Delinquency* 49, 1 (2012), 122–146. https://doi.org/10.1177/0022427810397950
[19] Peter Reuter and Davin O'Regan. 2017. Smuggling wildlife in the Americas: scale, methods, and links to other organised crimes. *Global Crime* 18, 2 (2017), 77–99.
[20] Janine E Robinson and Pablo Sinovas. 2018. Challenges of analyzing the global trade in CITES-listed wildlife. *Conservation Biology* 32, 5 (2018), 1203–1206.
[21] Gail Emilia Rosen and Katherine F Smith. 2010. Summarizing the evidence on the international trade in illegal wildlife. *EcoHealth* 7, 1 (2010), 24–32.
[22] Arunesh Sinha, Fei Fang, Bo An, Christopher Kiekintveld, and Milind Tambe. 2018. Stackelberg security games: Looking beyond a decade of success. IJCAI.
[23] William S Symes, Francesca L McGrath, Madhu Rao, and L Roman Carrasco. 2018. The gravity of wildlife trade. *Biological Conservation* 218 (2018), 268–276.
[24] Jason Tsai, Zhengyu Yin, Jun-young Kwak, David Kempe, Christopher Kiekintveld, and Milind Tambe. 2010. Urban security: Game-theoretic resource allocation in networked domains. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 24.
[25] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903* (2017).
[26] Ricardo Vilalta and Youssef Drissi. 2002. A perspective view and survey of meta-learning. *Artificial intelligence review* 18, 2 (2002), 77–95.
[27] Kai Wang, Andrew Perrault, Aditya Mate, and Milind Tambe. 2020. Scalable Game-Focused Learning of Adversary Models: Data-to-Decisions in Network Security Games.. In *AAMAS.* 1449–1457.
[28] Greg Warchol and Michael Harrington. 2016. Exploring the dynamics of South Africa's illegal abalone trade via routine activities theory. *Trends in Organized Crime* 19, 1 (2016), 21–41. https://doi.org/10.1007/s12117-016-9265-4
[29] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. 2020. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems* (2020).
[30] Tanya Wyatt, Daan van Uhm, and Angus Nurse. [n.d.]. Differentiating criminal networks in the illegal wildlife trade: organized, corporate and disorganized crime. *Trends in Organized Crime* ([n. d.]), 1–17.
[31] Lily Xu, Elizabeth Bondi, Fei Fang, Andrew Perrault, Kai Wang, and Milind Tambe. 2021. Dual-Mandate Patrols: Multi-Armed Bandits for Green Security. *AAAI* (2021).
[32] Lily Xu, Shahrzad Gholami, Sara Mc Carthy, Bistra Dilkina, Andrew Plumptre, Milind Tambe, Rohit Singh, Mustapha Nsubuga, Joshua Mabonga, Margaret Driciru, et al. 2020. Stay Ahead of Poachers: Illegal Wildlife Poaching Prediction and Patrol Planning Under Uncertainty with Field Test Evaluations (Short Version). In *2020 IEEE 36th International Conference on Data Engineering (ICDE).* IEEE, 1898–1901.
[33] Rong Yang, Fei Fang, Albert Xin Jiang, Karthik Rajagopal, Milind Tambe, and Rajiv T Maheswaran. 2012. Designing better strategies against human adversaries in network security games.. In *AAMAS.* 1299–1300.
[34] Bing Yu, Haoteng Yin, and Zhanxing Zhu. 2017. Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. *arXiv preprint arXiv:1709.04875* (2017).
[35] Jiani Zhang, Xingjian Shi, Junyuan Xie, Hao Ma, Irwin King, and Dit-Yan Yeung. 2018. Gaan: Gated attention networks for learning on large and spatiotemporal graphs. *arXiv preprint arXiv:1803.07294* (2018).
[36] Jie Zhou, Ganqu Cui, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. 2018. Graph neural networks: A review of methods and applications. *arXiv preprint arXiv:1812.08434* (2018).