

Robust Protection of Fisheries with COMPASS

William B. Haskell, Debarun Kar, Fei Fang and Milind Tambe
University of Southern California, Los Angeles, CA, 90089

Sam Cheung and Lt. Elizabeth Denicola
U.S. Coast Guard Research Development Center and Atlantic Area

Abstract

Fish stocks around the world are in danger from illegal fishing. In collaboration with the U.S. Coast Guard (USCG), we work to defend fisheries from illegal fisherman (henceforth called Lanchas) in the U.S. Gulf of Mexico. We have developed the COMPASS (Conservative Online Patrol ASSistant) system to design USCG patrols against the Lanchas. In this application, we face a population of Lanchas with heterogeneous behavior who fish frequently. We have some data about these Lanchas, but not enough to fit a statistical model. Previous security patrol assistants have focused on counter-terrorism in one-shot games where adversaries are assumed to be perfectly rational, and much less data about their behavior is available. COMPASS is novel because: (i) it emphasizes environmental crime; (ii) it is based on a repeated Stackelberg game; (iii) it allows for bounded rationality of the Lanchas and it offers a robust approach against the heterogeneity of the Lancha population; and (iv) it can learn from sparse Lancha data. We report the effectiveness of COMPASS in the Gulf in our numerical experiments based on real fish data. The COMPASS system is to be tested by USCG.

Introduction

Fish stocks are at risk of imminent collapse (see (Gillig, Griffin, and Ozuna Jr 2001; Baum and Myers 2004)). Illegal fishing greatly exacerbates the decline of fish stocks; it accounts for as much as 30% of the total catch in some major fisheries, and it is estimated that 11 - 26 million tons of fish were caught illegally worldwide each year (Petrossian 2012). As a specific example, the Gulf of Mexico is heavily exploited by illegal fisherman from Mexico, henceforth called Lanchas. The U.S. Coast Guard (USCG) wants to protect fisheries within the U.S. Exclusive Economic Zones from overfishing. The system reported in this paper is developed in collaboration with USCG with this aim.

Illegal fishing in the Gulf poses unique challenges to USCG. First, USCG can neither police nor observe the entire Gulf simultaneously. Second, it is assumed by USCG that Lanchas can observe their assets moving to and from their bases, and on the open water. Third, USCG does not

have perfect information about Lanchas making it difficult to forecast their future behavior. USCG detects many but not every Lancha incursion. Since we have limited data, it is hard to reliably fit a statistical model to observed Lancha behavior.

To respond to these challenges, we have developed the COMPASS (Conservative Online Patrol ASSistant) system. COMPASS is a decision aid for designing USCG patrols of the Gulf. This system is novel in four ways. First, COMPASS is a pioneering application of security game theory and robust optimization to environmental protection, while most research on security games emphasizes counter-terrorism (Tambe 2011; Paruchuri et al. 2008; Shieh et al. 2012; Conitzer 2012; Letchford and Vorobeychik 2011). Second, COMPASS models the interaction between Lanchas and USCG as a repeated Stackelberg game, while the literature on security games focuses on one-shot counter-terrorism games. Third, in contrast to earlier applications on security games which emphasize perfectly rational attackers, we adopt a bounded rationality model for Lancha behavior. However, we face an entire heterogeneous population of Lanchas and each could have a different behavioral model. COMPASS is novel because it is robust against this heterogeneity; the resulting patrols are conservative and protect USCG against uncertainty in Lancha type. Fourth, COMPASS combines robust optimization with learning to make use of available data to update its recommendations. In this way, it can still adapt to changes in Lancha behavior. COMPASS will be tested by USCG in the Gulf.

Domain

Our goal is to offer USCG a more efficient patrol strategy against Lanchas in the Gulf in the face of uncertainty about the true Lancha behavioral model. In the Gulf, the red snapper and shark populations are considered to be the most at risk from illegal fishing. Red snapper stocks are non-migratory, the location of these fisheries typically does not change over time. Sharks are migratory and spread throughout the Gulf. Empirically, it appears that the Lanchas have considerable knowledge about the distribution of these fish populations.

Lanchas enter U.S. waters frequently in small powerboats with crews of two to four, which are difficult to see against the water. Lanchas are either spotted directly, or their gear



Figure 1: (a) Coast Guard boat used to patrol Gulf of Mexico and (b) a captured Lancha boat

is detected. A captured Lancha and an USCG patrol asset is shown in Fig. 1. Throughout this paper, we assume that the Lanchas have knowledge of the distribution of USCG patrols.

Building COMPASS

We formally present the COMPASS system in this section. First we describe our game-theoretic framework and then we develop the details of COMPASS’s implementation.

Game-theoretic model

We model the interaction between USCG and Lanchas as a repeated Stackelberg game (Myerson 2013). USCG plays as the leader and generates patrols, then the Lanchas observe and react. This game is repeated because USCG can react to Lancha movement and design new patrols, which the Lanchas can then observe and use to improve their own strategy.

We discretize the Gulf into a grid of identical cells (targets) indexed by $\mathbb{T} = \{1, \dots, T\}$. We use real data on fish density and distance from the Mexican border to construct the payoff of each target to the Lanchas, R_t^a . R_t^a is directly proportional to the amount of fish in a grid cell but inversely proportional to the distance of that cell from their base to account for greater fuel cost. We also introduce a penalty term for Lanchas for the cost of being intercepted by USCG. P_t^a is the value of the Lancha vessel which is confiscated by USCG at target t if it is intercepted. If USCG protects target t , and the Lancha attacks target t , then USCG receives reward R_t^d from the Lancha penalty. If the Lancha attacks target t and it is not defended, then USCG loses P_t^d - the monetary value of the stolen fish. This game is zero sum. If the Lancha earns a high reward, they must have caught a lot of fish from near the shore (because most of the fish is near the shore). Thus, the USCG is penalized harder by missing a well known area of high fish density.

For r boat hours, USCG’s strategy can be represented by a probability distribution $x = (x_1, \dots, x_T)$ over targets in the set

$$\mathbb{X} = \left\{ x \in \mathbb{R}^T : \sum_{t \in \mathbb{T}} x_t = r, 0 \leq x_t \leq 1, \forall t \in \mathbb{T} \right\},$$

where x_t is the marginal probability of protecting target t and r is the total patrol boat hours. If the Lancha attacks

target t , then for a given defender strategy x , the defender’s expected utility is

$$U_t(x) = x_t R_t^d + (1 - x_t) P_t^d.$$

We are interested in randomized strategies because they make USCG more unpredictable. Each target t has a vector of attributes, in our case, the attribute vector will include the current USCG patrol density x_t , and the reward R_t^a and penalty P_t^a for that target. Actual patrols are generated from x taking into account time and distance constraints on USCG assets (Korzhyk, Conitzer, and Parr 2010; Jain et al. 2010).

We base our Lancha behavioral model on the Subjective Utility Quantal Response (SUQR) model from (Nguyen et al. 2013). SUQR is a stochastic choice model; it includes some randomness in adversary decision making which is not present in perfectly rational decision makers. SUQR gives us a realistic bounded rationality model of the Lanchas, since we cannot assume the Lanchas are perfectly rational. Further, (Nguyen et al. 2013) shows that SUQR outperforms other models like quantal response (McKelvey and Palfrey 1995) via extensive human subjects experiments in security games. In our case, using the data (x_t, P_t^a, R_t^a) available to the Lancha for target t , the probability that the Lancha will attack target t is

$$q_t(\omega | x) = \frac{e^{\omega_1 x_t + \omega_2 P_t^a + \omega_3 R_t^a}}{\sum_{v \in \mathbb{T}} e^{\omega_1 x_v + \omega_2 P_v^a + \omega_3 R_v^a}}.$$

Notice that $\sum_{t \in \mathbb{T}} q_t(\omega | x) = 1$. We see that the vector $\omega = (\omega_1, \omega_2, \omega_3)$ encodes all information about Lancha behavior, and each component of ω determines how much weight the Lancha gives the corresponding attribute in his decision making. In (Nguyen et al. 2013), ω is set to a fixed value to represent an attacker; there is no heterogeneity in attackers. Instead, we face an entire population of heterogeneous Lanchas, so we introduce a set $\Omega \subset \mathbb{R}^3$ to represent the range of all possible ω that represent real Lanchas. We allow a *separate* value of ω for each Lancha incursion. We will refer to ω as a specific Lancha type.

Given that we have heterogeneous Lanchas, rather than just one, we are dealing with a population of ω . A natural idea is to assume that there is a prior distribution F over Ω . Then, the stochastic optimization problem

$$\max_{x \in \mathbb{X}} \int_{\Omega} \left[\sum_{t \in \mathbb{T}} U_t(x) q_t(\omega | x) \right] F(d\omega) \quad (1)$$

can be solved. Problem (1) maximizes the expected utility for USCG, where expectation is taken over Lancha types. Furthermore, if Lancha data is available, then Bayesian updates can be performed on F .

However, there are several problems with the stochastic optimization approach in Problem (1). First, we cannot justify the general assumption that there is a probability distribution F over Lancha types. There is just not enough data to make a reasonable hypothesis about the distribution of ω for the Lancha population. Second, when observations of ω are available we can perform Bayesian updating in Problem (1)

to improve our estimate of the true probability distribution on Ω . However, this type of Bayesian updating requires a lot of data even if a prior distribution F is available. Even when a lot of data is available, it may take time to gather these data and Bayesian updating would not be effective while that data gathering takes place. Runtime presents itself as a third issue, since solving Problem (1) is computationally expensive.

Robust optimization

Robust optimization offers remedies for the shortcomings in Problem (1). It does not require a distribution F over Lancha population parameters. Also, robust optimization will work with very small data sets - even a single observation. We continue to use the SUQR model from (Nguyen et al. 2013), since we do not rely on the assumption of perfectly rational attackers as in earlier security games. We now enhance the SUQR model by hedging against the heterogeneity of the Lancha population with robust optimization.

We start with an uncertainty set $\hat{\Omega} \subset \Omega$ to capture the possible range of reasonable Lancha behavior parameters. $\hat{\Omega}$ can be initialized based on domain expertise and human behavior experiments. As new observations come in, COMPASS updates $\hat{\Omega}$ as described in the next section.

Given an uncertainty set $\hat{\Omega}$, we solve the initial robust optimization problem

$$\max_{x \in \mathbb{X}} \min_{\omega \in \hat{\Omega}} \sum_{t \in \mathbb{T}} U_t(x) q_t(\omega | x), \quad (2)$$

to get a patrol for USCG. We continue to assume that individual Lanchas follow an SUQR model. However, we do not know the distribution of the SUQR parameters ω over this population. Problem (2) is distinct from Problem (1). No probability distribution is assumed and no Bayesian updating is performed in Problem (2). Additionally, the objective of Problem (1) is the USCG expected utility over SUQR parameters with respect to F . In contrast, the objective of Problem (2) is the worst-case expected utility over all allowable SUQR models for the Lancha. Since we do not know the distribution of ω , it is safest for us to hedge against all possible Lancha types.

Problem (2) is a nonlinear, nonconvex, nondifferentiable optimization problem. The function $\sum_{t \in \mathbb{T}} U_t(x) q_t(\omega | x)$ is continuous and differentiable in x for any fixed ω , but it is not convex in x . For easier implementation, we transform Problem (2) into the constrained problem

$$\max_{x \in \mathbb{X}, s} \left\{ s : s \leq \sum_{t \in \mathbb{T}} U_t(x) q_t(\omega | x), \forall \omega \in \hat{\Omega} \right\}, \quad (3)$$

by introducing a dummy variable $s \in \mathbb{R}$. This transformation replaces the nonsmooth objective in the original Problem (2) with a set of smooth constraints, so that we eliminate the nondifferentiability to make numerical algorithms work better. Specifically, we solve Problem (2) in MATLAB with the function *fmincon*.

Learning

This section explains how COMPASS uses data. Learning is a core part of COMPASS, and any available data can improve our online patrol generation. The main intuition is that we will use data to construct the uncertainty set $\hat{\Omega}$ for Lancha types that was taken as input in Problem (2). In this way, we keep the conservatism of robust optimization but include the optimism that we can adapt to new observations. This idea is based on (Bandi and Bertsimas 2012), where it is used to construct uncertainty sets for stochastic analysis. This blend of robustness and learning has not yet appeared in the security game literature, and it is a promising approach to patrol scheduling in security games where some data is available. Unlike other estimation-based approaches, our approach does not require strong statistical assumptions on the underlying data (like normality).

We will assume that each observation comes from an independently distributed Lancha attack, with a possibly different SUQR parameter ω from our population of parameters. We do not have the ability to know which the repeat criminals are, so we have to assume that all incursions are independent and each corresponds to a different SUQR parameter. Let $N^{(k)}$ be the total number of crimes committed in round k , and let $J^{(k)} = \sum_{i=1}^k N^{(i)}$ be the running total number of crimes by the end of round k . Suppose target t is attacked during round k while the USCG patrol density is x_k , then the maximum likelihood equation (MLE) for the type of the Lancha, i.e. the parameter ω , who committed this crime is

$$\max_{\omega \in \Omega} \log(q_t(\omega | x_k)). \quad (4)$$

To explain, equation (4) finds the Lancha type ω which maximizes the likelihood of an attack on the observed target t . Note that equation (4) only assumes the SUQR model, it is not based on any other hidden assumptions on the data. By solving the MLE equation (4) for each observed incursion, we build a collection

$$\Xi^{(k)} = \left\{ \omega_j : j = 1, \dots, J^{(k)} \right\}$$

of estimated Lancha SUQR parameters available at the end of round k . The set $\Xi^{(k)}$ is effectively a sample of estimated Lancha SUQR parameters, we will use $\Xi^{(k)}$ to construct $\hat{\Omega}$.

For clarity, $\Xi^{(k)}$ is not known until the **end** of round k . At the beginning of round k , USCG only knows $\Xi^{(k-1)}$. We denote the resulting data-driven uncertainty set for round k as

$$\hat{\Omega} \left(\Xi^{(k-1)} \right)$$

to emphasize the dependence on the data $\Xi^{(k-1)}$. At the beginning of round k , we then want to solve the resulting robust optimization problem

$$\max_{x \in \mathbb{X}} \min_{\omega \in \hat{\Omega}(\Xi^{(k-1)})} \sum_{t \in \mathbb{T}} U_t(x) q_t(\omega | x), \quad (5)$$

which can be solved using the same algorithm as for Problem (2). Compare Problems (2) and (5). In Problem (2), the uncertainty set $\hat{\Omega}$ is initialized and no new data has yet been

collected. In Problem (5), the uncertainty set $\widehat{\Omega}(\Xi^{(k-1)})$ includes new Lancha observations over the last $k - 1$ rounds. Different data sets give rise to different instances of the robust Problem (5). The only remaining issue is to construct the data-driven uncertainty set $\widehat{\Omega}(\Xi^{(k-1)})$.

We define our uncertainty sets in the following way: In round k , we consider

$$\begin{aligned} & \widehat{\Omega}(\Xi^{(k-1)}) \\ &= \text{vertex} \left\{ \omega_j : j = J^{(k-1)} - m(k), \dots, J^{(k-1)} \right\}. \end{aligned}$$

where $\text{vertex}(\cdot)$ returns the vertices of the convex hull of a set and $m(k)$ is the number of past observations that we keep. If $m(k) = J^{(k-1)} - 1$, then we use all Lancha observations. We can easily make $\widehat{\Omega}(\Xi^{(k-1)})$ less conservative by having a shorter memory. In this fashion, we do not become progressively more conservative but track recent Lancha activity.

Unlike Problem (2), where $\widehat{\Omega}$ is initialized by the user, perhaps based on historical data, the uncertainty set $\widehat{\Omega}(\Xi^{(k-1)})$ is data-driven and is updated online each time a new Lancha observation is made. In fact there are many ways to construct $\widehat{\Omega}(\Xi^{(k-1)})$.

Experiments and Analysis

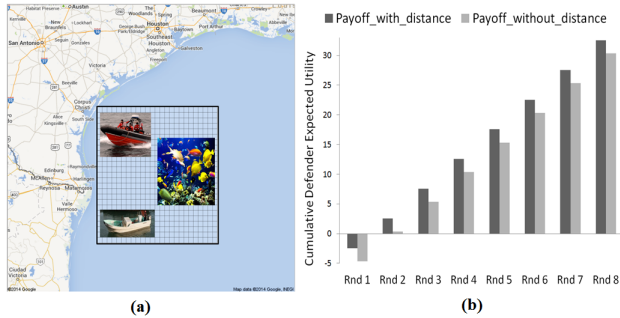


Figure 2: (a) Sample Simulation Area in Gulf of Mexico; (b) Cumulative Expected Utility for COMPASS with and without the distance factor being included in the payoffs

In this section, we evaluate COMPASS on real fish and USCG Lancha data. The fish data was taken from NOAA’s Coastal Ecosystem Maps¹; for security purposes we had to sanitize the Lancha data obtained from USCG to share these results. We focus on part of the Gulf of Mexico and discretize it into a 25*25 grid with 625 cells, 3 miles wide and 4.2 miles long. This area is shown in Fig. 2, and each round corresponds to 3 months. For these experiments, the Lanchas are uniformly penalized across targets if they are caught by USCG. We will return to Fig. 2(b) later.

We evaluate COMPASS and compare it against other algorithms. Figs. 3(a), 3(b) and 3(c) depict how the cumulative expected utility of the defender (shown on the Y-axis)

increases over rounds (shown on the X-axis) for 10, 30 and 50 patrol boat hours, respectively. So, we calculate the worst case expected utility for the defender over all Lancha types for each round and this gets cumulated over the rounds. We see that COMPASS improves based on an average of about 30 Lancha observations per round. This performance is consistent even when varying the number of resources. We note the superiority of COMPASS over two other approaches: first, we consider a baseline Maximin that does not use a behavioral model; second, we compare against SUQR. For SUQR, we assume that the weights $\omega = (\omega_1, \omega_2, \omega_3)$ are the same for all Lanchas, we use the weights derived from human subject experiments in (Nguyen et al. 2013). The SUQR model performs the worst because it fails to exploit the fact that different Lanchas may use different weights for different features, and because it does not update over rounds. On the other hand, COMPASS successfully captures a diverse population of Lanchas and learns from data. Even though the performance of both Maximin and SUQR improve as the number of boat hours increases, they are consistently outperformed by COMPASS. Also note that COMPASS utilizes resources more efficiently as compared to Maximin and SUQR, which are both sensitive to the number of resources.

The heat maps in Figs. 4(a) and 4(b) reveal the coverage probabilities generated by COMPASS for rounds 1 and 7, lighter colors correspond to higher coverage probabilities. These figures clearly show that COMPASS’s robust approach is capable of adapting to data. Note that the grid to the bottom right corner of Figs. 4(a) and 4(b) shows the heat map for the entire 25*25 grid, while the heat map of a zoomed in 3*3 grid is shown at the left upper corner of the figures. The 3*3 grid corresponds to rows 5 to 8 and columns 1 to 3 in the 25*25 grid. Similar notation is followed for Fig. 4(c), which shows a heat map of the strategy generated by Maximin. In contrast to the heatmaps of COMPASS which change over rounds, Maximin generates only one fixed strategy for every round. We also experiment with COMPASS to study the impact of distance from the shore on the defender strategy. Fig. 2(b) reports the cumulative defender expected utility for COMPASS for two cases: (i) when the distance from shore is incorporated into the rewards and (ii) when it is not. The results suggest that traveling distance is an important factor to consider while constructing the payoffs.

Related Work

Our work synthesizes the literature on fishery protection, security games, and robust optimization. First, the problem of protecting fisheries has already received attention. (Pauly et al. 2002) discusses the urgent need for protection of fisheries to maintain sustainability. In (Gallic and Cox 2006; Stokke 2009), an economic analysis of illegal fishing is conducted and some economic remedies are proposed. Our work differs from these references because it offers a tactical, proactive solution to illegal fishing rather than a strategic, policy based solution. (Lobo 2005) solves for one-dimensional ocean patrols using self-organizing maps and (Millar and Russell 2012) uses an integer programming formulation to design fishery patrol routes. Both (Lobo 2005) and (Millar and Russell 2012) are based on the classic trav-

¹<http://service.nccdc.noaa.gov/website/CHP/viewer.htm>

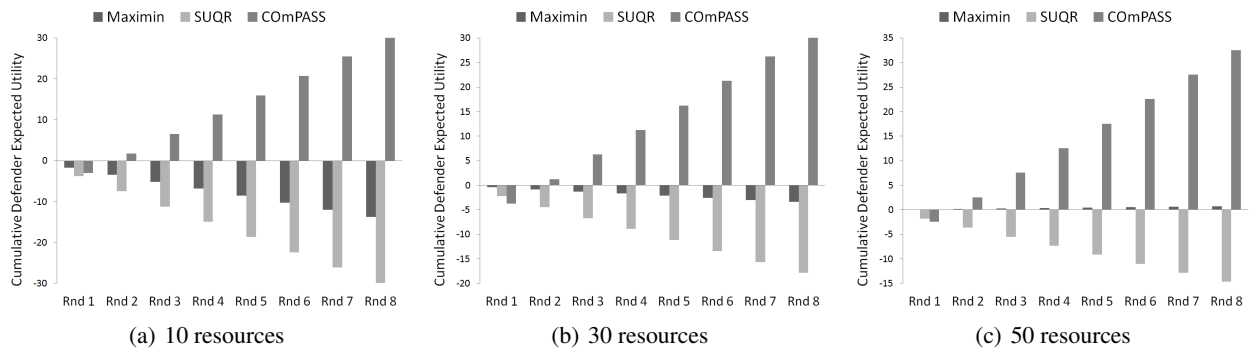


Figure 3: Cumulative Defender Expected Utilities for different number of resources over rounds

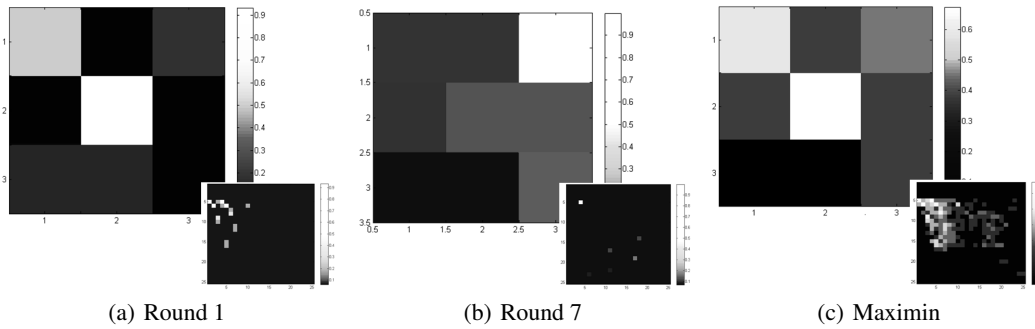


Figure 4: Heat Maps showing defender coverage probabilities for all targets generated by (a, b) ComPASS with 50 resources and (c) Maximin

eling salesman problem. Our approach differs from (Lobo 2005) in two ways: first, we take a game-theoretic approach; second, we account for uncertainty in the behavior and preferences of the adversary to get robust patrols. (Yang et al. 2014) studies a wildlife security game against a population of heterogeneous poachers. The SUQR model also appears in this work, and it is assumed that the SUQR parameters for the poachers are normally distributed. ComPASS does not make any assumptions on the distribution of SUQR parameters.

Game theory in security is a highly developed field, especially in counter-terrorism applications (Tambe 2011). The Stackelberg game model has been emphasized in this body of literature. (Conitzer 2012) surveys the state of the art in algorithmic security game theory, in particular for Stackelberg games; though these algorithms do not apply to our repeated setting. Furthermore, Stackelberg games in security often appear in conjunction with a behavioral model for the adversary. (Marecki, Tesauro, and Segal 2012) considers repeated Stackelberg Bayesian games, where the prior distribution is over follower preferences and the follower is assumed to play a myopic best response. We depart from this approach because we do not assume a prior over adversary preferences, and we work in the specific fisheries domain.

Robust optimization, a tool for decision-making under uncertainty, is particularly pertinent to us. The usual strategy is to optimize a performance measure over the worst-case of scenarios, thus ensuring a conservative decision. See (Ben-

Tal, El Ghaoui, and Nemirovski 2009) for a detailed survey of robust optimization and see (Bandi and Bertsimas 2012) for the application of robustness to stochastic systems analysis. Robustness is especially relevant to multiple agent settings, and robust solution concepts have been developed for games. (Aghassi and Bertsimas 2006) develop a solution concept for robust Nash games, and (Kardeş, Ordóñez, and Hall 2011) develop a solution concept for robust Markov games. We differ from the usual robust approach because: (i) we are focused on game theory; (ii) we are solving a sequence of robust optimization problems; and (iii) our robust formulation makes use of data.

Summary

ComPASS is a novel application of security game theory and robust optimization that improves USCG’s ability to interdict Lanchas in the Gulf of Mexico. To summarize, ComPASS has four main novel features. First, ComPASS expands the application of security game theory and robust optimization to the critical task of protecting fisheries. Second, ComPASS contributes a model of repeated Stackelberg games to the security literature, which traditionally emphasizes one-shot games. Third, we develop a robust variant of SUQR that protects against uncertainty in the Lancha population. This robustness is the key to generating defensive patrols. Finally, ComPASS offers a way to dynamically update our uncertainty set online as we gather new observations

so that we can adapt to changing Lancha behavior. USCG is studying our recommendations to inform their patrols of sea and air assets. Whenever our patrols are implemented, USCG will take detailed observations and return this information to us. Their domain experts are able to effectively comment on our patrol feasibility and impact. We conclude by emphasizing that our method has wide applicability to other security games beyond fishery protection where there is repeated interaction between the population of criminals and law enforcement. The SUQR framework is a general human behavioral model for making choices, it is not specific to the fisheries domain.

Disclaimer

The views expressed in this article are those of the author and DO NOT necessarily reflect the views of the USCG, DHS, or the U.S. Federal Government or any Authorized Representative thereof. The U.S. Government does not approve/ endorse the contents of this article, therefore the article shall not be used for advertising or endorsement purposes. The U.S. Government assumes NO liability for its contents or use thereof. Moreover, the information contained within this article is NOT OFFICIAL U.S. Government policy and cannot be construed as official in any way. Also, the reference herein to any specific commercial entity, products, process, or service by trade name, trademark, manufacturer, or otherwise, does NOT necessarily constitute or imply its approval, endorsement, or recommendation by the U.S. Government.

References

Aghassi, M., and Bertsimas, D. 2006. Robust game theory. *Mathematical Programming Series B* 107:231–273.

Bandi, C., and Bertsimas, D. 2012. Tractable stochastic analysis in high dimensions via robust optimization. *Mathematical programming* 134(1):23–70.

Baum, J. K., and Myers, R. A. 2004. Shifting baselines and the decline of pelagic sharks in the gulf of mexico. *Ecology Letters* 7(2):135–145.

Ben-Tal, A.; El Ghaoui, L.; and Nemirovski, A. 2009. *Robust optimization*. Princeton University Press.

Conitzer, V. 2012. Computing game-theoretic solutions and applications to security. In *AAAI*, 2106–2112.

Gallic, B. L., and Cox, A. 2006. An economic analysis of illegal, unreported and unregulated (iuu) fishing: Key drivers and possible solutions. *Marine Policy* 30(6):689–695.

Gibbons, R. 1992. *Game theory for applied economists*. Princeton University Press.

Gillig, D.; Griffin, W. L.; and Ozuna Jr, T. 2001. A bio-economic assessment of gulf of mexico red snapper management policies. *Transactions of the American Fisheries Society* 130(1):117–129.

Jain, M.; Kardeş, E.; Kiekintveld, C.; Ordonez, F.; and Tambe, M. 2010. Optimal defender allocation for massive security games: A branch and price approach. In *AA-MAS 2010 Workshop on Optimisation in Multi-Agent Systems (OptMas)*.

Kardeş, E.; Ordóñez, F.; and Hall, R. W. 2011. Discounted robust stochastic games and an application to queueing control. *Operations research* 59(2):365–382.

Korzhyk, D.; Conitzer, V.; and Parr, R. 2010. Complexity of computing optimal stackelberg strategies in security resource allocation games. In *AAAI*.

Letchford, J., and Vorobeychik, Y. 2011. Computing randomized security strategies in networked domains. In *AARM Workshop In AAAI*.

Lobo, V. 2005. One dimensional self-organizing maps to optimize marine patrol activities. In *Oceans 2005-Europe*, volume 1, 569–572. IEEE.

Marecki, J.; Tesauro, G.; and Segal, R. 2012. Playing repeated stackelberg games with unknown opponents. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, 821–828. International Foundation for Autonomous Agents and Multiagent Systems.

McKelvey, R. D., and Palfrey, T. R. 1995. Quantal response equilibria for normal form games. *Games and Economic Behavior* 2:6–38.

Millar, H. H., and Russell, S. N. 2012. A model for fisheries patrol dispatch in the canadian atlantic offshore fishery. *Ocean and Coastal Management* 60(0):48 – 55.

Myerson, R. B. 2013. *Game theory: analysis of conflict*. Harvard university press.

Nguyen, T. H.; Yang, R.; Azaria, A.; Kraus, S.; and Tambe, M. 2013. Analyzing the effectiveness of adversary modeling in security games. In *Conf. on Artificial Intelligence (AAAI)*.

Paruchuri, P.; Pearce, J. P.; Marecki, J.; Tambe, M.; Ordonez, F.; and Kraus, S. 2008. Playing games for security: An efficient exact algorithm for solving bayesian stackelberg games. In *AAMAS*.

Pauly, D.; Christensen, V.; Guénette, S.; Pitcher, T. J.; Sumaila, U. R.; Walters, C. J.; Watson, R.; and Zeller, D. 2002. Towards sustainability in world fisheries. *Nature* 418(6898):689–695.

Petrossian, G. A. 2012. *The decision to engage in illegal fishing: An examination of situational factors in 54 countries*. Ph.D. Dissertation, Rutgers University-Graduate School-Newark.

Shieh, E.; An, B.; Yang, R.; Tambe, M.; Baldwin, C.; Di-Renzo, J.; Maule, B.; and Meyer, G. 2012. Protect: A deployed game theoretic system to protect the ports of the united states. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Stokke, O. S. 2009. Trade measures and the combat of iuu fishing: Institutional interplay and effective governance in the northeast atlantic. *Marine Policy* 33(2):339–349.

Tambe, M. 2011. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. New York, NY: Cambridge University Press.

Yang, R.; Ford, B.; Tambe, M.; and Lemieux, A. 2014. Adaptive resource allocation for wildlife protection against illegal poachers. *AAMAS*.