

From the Lab to the Classroom and Beyond: Extending a Game-Based Research Platform for Teaching AI to Diverse Audiences

Nicole Sintov¹, Debarun Kar¹, Thanh Nguyen¹, Fei Fang¹,
Kevin Hoffman², Arnaud Lyet³, Milind Tambe¹

¹{sintov,dkar,thanhhng,feifang,tambe}@usc.edu, University of Southern California, USA

²khoff97@gmail.com, Aspire Public Schools, USA

³Arnaud.Lyet@wwfus.org, World Wildlife Fund, USA

Abstract

Recent years have seen increasing interest in AI from outside the AI community. This is partly due to applications based on AI that have been used in real-world domains, for example, the successful deployment of game theory-based decision aids in security domains. This paper describes our teaching approach for introducing the AI concepts underlying security games to diverse audiences. We adapted a game-based research platform that served as a testbed for recent research advances in computational game theory into a set of interactive role-playing games. We guided learners in playing these games as part of our teaching strategy, which also included didactic instruction and interactive exercises on broader AI topics. We describe our experience in applying this teaching approach to diverse audiences, including students of an urban public high school, university undergraduates, and security domain experts who protect wildlife. We evaluate our approach based on results from the games and participant surveys.

Introduction

In recent years, the field of security games (Tambe 2011), a subfield of AI, has drawn increasing attention from outside the artificial intelligence (AI) community. In particular, game theory-based decision aids for protecting critical infrastructure such as airports (Pita et al. 2008) and ports (Shieh et al. 2012) have been successfully deployed, making real-world impacts and resulting in fundamental changes to security operations for select organizations.

As a subfield of computational game theory, security games model the strategic interaction between two players, a defender and an adversary. The defender allocates security resources to protect a set of critical targets of varying importance. The defender's strategy is a probability distribution over all possible ways of allocating the resources to targets. The adversary conducts surveillance on the defender's strategy before selecting a target to attack. The players' actions lead to payoff values, and the defender's performance is evaluated by her expected utility, i.e., her payoff in expectation when a defender strategy is played. The defender's

goal is to find the optimal strategy so as to maximize her expected utility, knowing she faces an adaptive adversary who will respond to any deployed strategy.

The existing and potential applications of this work have created the need to introduce the AI concepts underlying security games to individuals with limited AI backgrounds, and not only to students, but also to audiences beyond the classroom. This has become particularly relevant following recent advances in the sub-field of green security games, which use game-theoretic strategies to protect forests (Johnson, Fang, and Tambe 2012), fisheries (Haskell et al. 2014) and wildlife (Fang, Stone, and Tambe 2015; Kar et al. 2015). In fact, wildlife protection is a global concern. Many species, including tigers and rhinos, are in danger of extinction due to poaching. Although park rangers have been conducting patrols over conservation areas to combat this illegal activity, security resources are often limited in vast park areas. Manually generating patrol schedules can require considerable effort from wildlife staff, and such manual plans also can be predictable, allowing poachers to exploit patrol schedules. Our security game-based solutions combine different AI subfields, including game theory, optimization, and machine learning, to help rangers in *automatically* generating *randomized* patrol strategies that account for models of poachers' behaviors.

Teaching these concepts and techniques to those with limited AI backgrounds can be challenging. In particular, as described above, the defender's quantitative decision, which is a probability distribution over his possible actions, is made based on reasoning about both players' payoffs associated with their actions (Yin et al. 2010). This requires reviewing general AI concepts, probability, and optimization. While general AI concepts and techniques have been made accessible to undergraduate students (Stern and Sterling 1996; Parsons and Sklar 2004; Wollowski 2014), no prior work describes effective methods for teaching security games to audiences in the classroom or beyond. Bringing these concepts to audiences beyond the classroom is particularly important: it is necessary for decision-makers and those who may consider using AI-based decision aids in the field to understand the underlying theoretical framework, in part to increase adoption of these emerging technologies.

One method that has been effective in teaching general AI in classrooms is the use of games (Wong, Zink, and Koenig

2010; DeNero and Klein 2010). However, little evidence speaks to the efficacy of using security games to pique interest in or teach about AI in general. Further, no data exists on teaching these concepts beyond the classroom.

This paper describes our novel teaching approach that leveraged a game-based research platform, and moved it to the classroom and beyond. Our computer-based game was initially developed and used to gather data via human subjects experiments in a study of green security games (Kar et al. 2015). We modified the game to increase suitability as a hands-on learning tool. We describe this approach and our experiences introducing general AI and security game concepts to several audiences, including: (1) students of an urban public high school; (2) undergraduate students at a large private university; and (3) law enforcement officers and rangers who work in wildlife protection in Indonesia. In teaching each of these audiences, we delivered didactic instruction on AI and game theory, followed by guiding learners in playing the role-playing games.

We evaluate our teaching approach in two ways. First, we show how participants played the games and the effectiveness of their defender strategies, represented as defender utility, i.e., the defender's overall payoff obtained when peers play as poachers. Second, we assess participants' perceived value of our approach using surveys. We make recommendations for teaching similar topics to audiences with limited AI backgrounds.

Background and Role-Playing Games

As described above, green security games involve the development and application of game-theoretic strategies to protect forests (Johnson, Fang, and Tambe 2012), fisheries (Haskell et al. 2014) and wildlife (Fang, Stone, and Tambe 2015; Kar et al. 2015). A game-theory based decision-aid called PAWS (Yang et al. 2014) (Protection Assistant for Wildlife Security) was developed in part based on a study of green security games for the purpose of protecting wildlife from poaching. The study used an online computer game (Figure 1) to understand human decision making through a set of human subjects experiments (Kar et al. 2015). Our role-playing games were developed based on this game.

The game begins with a background story in the domain of wildlife protection. More specifically, the game simulates a real-world poaching scenario in Queen Elizabeth National Park (QENP) in Uganda. An image of the park is depicted and is discretized into a 5 by 5 grid. Animals (hippos) are shown in varying densities across the grid. There are two players: (1) the law enforcement officer (hereafter referred to as defender or ranger) and (2) the poacher. The defender moves first, deciding a patrol strategy, which is a randomized patrol schedule for park rangers. This is represented as the probability a ranger will be in each grid cell. Then the poacher can walk through the park to see the relative odds of being successful in a poaching expedition, and decide to poach / place a snare in a given area (grid cell). Then a sample of the rangers' actual location is randomly generated from the patrol strategy. If a ranger is in the same grid cell as the poacher, the poacher will get caught and lose money, otherwise the poacher wins and gains money. The computer



Figure 1: Online computer game interface (Kar et al. 2015).

generates the defender strategy according to a game-theory based algorithm and the participants play as poachers. In the original study, the game was played over multiple rounds. After each round, data from participants' actions were collected, based on which a new defender strategy was generated for use in the next round.

Players in the original game experiment always played as poachers. They were first primed with a background story about the hardships of a poacher's life, then instructed to look throughout the park to place a snare to catch a hippopotamus. We adapted this for our teaching approach by creating a set of paired role-playing games and the opportunity for learners to play as both poachers and rangers. Specifically, we created board games mimicking the computer-based game. Board games used the same background story and 5 by 5 grid as the online game. However, instead of using the QENP map and images of hippos to show the relative value of different areas, the board games used movable figures to represent animal distributions. This did not require any facilities such as computers, making it easily scalable to other educational settings. Also, in board games, learners took turns playing the roles of ranger and poacher, moving figures of rangers across the board to generate the defender strategies, and moving poacher figures to represent poacher decision-making. A subset of our audiences also played an adapted version of the computer-based game as poachers.

High School Students

Background An AI unit was delivered to a group of high school students as the last unit of study in a year-long engineering elective course. The structure of the unit was designed based on Dr. Milind Tambe's undergraduate-level course at the University of Southern California entitled CS499: Artificial Intelligence and Science Fiction (Tambe, Balsamo, and Bowring 2008). The high school unit was developed during summer 2014 in Dr. Tambe's Teamcore research group, funded by the National Science Foundation's Research Experience for Teachers (RET) program called ACCESS 4Teachers. The primary objective of this unit was to introduce state-of-the-art concepts in AI to high school students. In addition, basic concepts in game theory were introduced. Students also participated in the interactive role-playing games including: (i) playing the board games as

rangers, designing defender strategies; and (ii) playing the board game and computer game as poachers against the defender strategies developed by other students in the class and also the AI-generated algorithms.

Participants The 30 students who participated in the elective course were juniors and seniors at an urban public high school located in Los Angeles. All students were of Hispanic or African-American origin, resided in South Los Angeles, and the majority qualified for free or reduced price lunch.

Overall Course Content A key tenet of this unit was that AI concepts can be made accessible to anyone in the right context. We also placed an emphasis thinking about problems from a quantitative perspective, and considering how human-like qualities (emotions, risk-aversion, etc.) could be quantified to enable computers to act in an intelligent fashion. Students read Asimov and watched *Stark Trek*, and debated the meaning of intelligence based on these stories. They were then taught about how quantitative decisions could be made. The teacher delivered lectures on unit content, and skills were reinforced by students completing worksheets and applied problems. More pop-culture examples were given, ranging from an exploration of the California Lottery to the popular game show *Deal or No Deal?*. For many students, a full grasp of the material required a review of fraction operations and probability, as well as an introduction to the concept of expected value.

Final Project The final project for the unit used the role-playing board game. To reinforce the idea of quantitative decision-making, students were tasked with designing their own defender strategies using short formulas to allocate a limited number of ranger-hours to the 25 grid cells in the game. Students worked independently or in groups to complete the strategies (in this case, using Google spreadsheets). These spreadsheets were then used to generate twelve distinct games based on the student-designed defense strategies. Taking on the role of poacher, students played the games against defender strategies generated by other groups. Finally, students reviewed the results of their strategies, made adjustments, and presented their work to explain where initial strategies were particularly successful or unsuccessful.

Results of Role-Playing Game Different groups employed different methods in designing defender strategies. While some chose to concentrate ranger patrols in areas dense with animals, often associated with a high probability of failure, others developed strategies in which the expected value for poachers was nearly zero. Fig. 2 shows the expected defender utilities obtained by each of the 12 student groups based on the attacks conducted on the other teams' defender strategies. The team with the highest expected utility generated a strategy that not only considered the animal densities, but also the distance from the poacher's starting location, placing lower coverage in cells farther away from the poacher's starting location. The team with the lowest expected utility placed maximum coverage (100%) on the highest animal density cell and divided remaining resources (ranger-hours) uniformly across all remaining targets, thus




Figure 2: Defender utilities for each high school student group.

ignoring important factors like animal distribution and distance.

Feedback A total of 13 (7 males, 5 females, 1 declining to state) out of 30 students responded to a survey that assessed their experiences in the unit. Questions assessed perceived educational value of the unit, likelihood of recommending the unit to others, and the unit's impact on their overall interest in AI. Responses were provided on likert scales (e.g., ranging from 1=strongly disagree to 7=strongly agree). Open-ended questions assessed general likes and dislikes. The mean age of respondents was 18.1 years (SD=0.8). More than 90% responded that they would recommend the activity to other high school students. Additionally, more than 70% agreed (somewhat or more) that the activity increased their interest in AI, and 93% agreed (somewhat or more) that the activity was a valuable learning experience. Open-ended responses indicated that respondents particularly enjoyed the interactive game aspect of the course. The least enjoyable aspects of the course were those that students felt were tedious or repetitive, such as calculating probabilities.

University Students

Background AI and Security Games were introduced to a class of University of Southern California (USC) Freshmen as a 2-week unit in their Freshman Academy course in fall 2015, which is an introductory engineering course. Similar to the high school course described in the previous section, the 2-week AI unit portion of the course was also designed based on Dr. Tambe's seminar at USC entitled CS499: Artificial Intelligence and Science Fiction. In addition to the introduction of basic concepts in AI, as part of the AI unit, the students also participated in interactive activities including: (i) playing board games as rangers to design strategies to improve patrolling effectiveness; (ii) playing board games as poachers to measure the effectiveness of strategies developed by peers and also the AI algorithms.

Participants The 30 students who took part in the AI unit were all USC freshmen majoring in engineering. As mentioned, the unit was part of a 2-unit course aimed at introducing students to ongoing research at USC across various engineering disciplines.

Overall Course Content The unit consisted of didactic instruction and discussion of basic concepts in AI and game theory, including applications to solving real-world problems. University professors and CEOs of technical startups




Figure 3: Defender utilities obtained by each group

facilitated discussion around the use of AI in security applications, painting a comprehensive picture of the various ways in which AI can influence day-to-day life. To understand different concepts such as the notion of payoffs in a game context, the classic Prisoner's dilemma problem in game theory was discussed. This provided foundational knowledge to facilitate learning about the payoffs in the role-playing games (i.e., the animal densities and penalties).

Similar to the final project for high school students in the previous section, following the didactic and discussion-based elements of the unit, students played the role-playing board game that mimicked the computer-based game. Here, students first played as rangers. The class was divided into seven groups, each of which designed its own defender strategy on a game board. Some groups chose to allocate ranger coverage in proportion to the number of animals, while others placed highest probability of capture at the highest animal density region and uniform everywhere else; some others developed strategies in which the expected value for poachers was nearly zero across all of the regions of the park. Each group's strategy was then shown to the other groups, who played the game in the poacher role against their peers' defender strategies.

The following week, the results of the board games were presented. Basic game-theoretic concepts such as Maximin were then explained. This led students to focus on subsets of information in decision-making: for instance, in the case of Maximin, when only information about the payoffs in the game is available, how can the most conservative strategy be developed? This was an important concept to help teams learn how to improve their strategies.

Results of Role-Playing Game Complementing the high school students' results (Fig. 2), the resulting defender utilities for each university student group playing the board games is shown in Fig. 3. The team with the lowest defender utility (G3) placed very low coverage (< 0.40) in the highest animal density cell and as expected, all the other teams attacked that cell. Similar results were obtained for group G7 which placed a coverage of 0.50 on the highest animal density target. Other teams performed reasonably well but none performed better than Maximin, which would have resulted in a defender expected utility of -2.17.

Feedback A total of 24 (7 males and 17 females) out of 30 students responded to a survey that assessed their experiences in the unit. Questions mirrored those administered to the high school students. The mean age of the sample was 18.22 years. More than 65% of students responded that they would recommend the activity to other university students.



(a) Computer games

(b) Board games

Figure 4: Participants played Role Playing Games

Additionally, more than 69% of students noted that the activity increased their interest in AI at least somewhat, and more than 80% agreed (somewhat or more) that the activity was a valuable learning experience. Qualitative data indicated that respondents particularly enjoyed the interactive aspect of the unit. The least enjoyable aspects of the unit were cases where students in a team couldn't agree on a particular strategy.

Workshop on Wildlife Protection in Indonesia

Overall Course Content A 3-day workshop was developed in collaboration with the World Wildlife Fund (WWF) to demonstrate the value of AI-based solutions for security (i.e., PAWS software), to security experts who protect wildlife. The workshop was held in Bandar Lampung, Sumatra, Indonesia in May 2015. We sought to teach how AI algorithms fed with partial information can generate patrol strategies that can outperform those generated by field experts with extensive knowledge of the system. We used the paired role-playing games to teach about these concepts and real-world applications that leverage them. In addition, we delivered didactic tutorials on theory and applications of agent-based modeling, game theoretic concepts, and security games. Learners also played the paired role-playing games, playing the computer-based games and board games as poachers and rangers. They also discussed in groups various challenges faced in wildlife protection and solutions for those challenges. Finally, they generated and shared ideas for improving patrolling effectiveness. These interactive exercises provided learners with a new lens for understanding poachers' behaviors and weaknesses of manual patrolling strategies, as well as introducing the methodology and advantages of PAWS game-theoretic solutions.

Participants A total of 28 participants (26 males and 2 females) attended the PAWS workshop representing the five primary groups (either Government or NGOs) involved in protecting wildlife in Bukit Barisan Selatan (BBS) and Tesso Nilo (TN) national parks in Indonesia: Indonesian National Park Service, WWF, Wildlife Conservation Society (WCS), Indonesian Rhino Foundation (YABI), and Prosecution Officers from the Court. These individuals were park rangers and law enforcement officers with a great deal of experience in wildlife protection and domain expertise in wildlife crime and protection. The mean age was 35.0 years ($SD = 7.5$), and mean years of formal schooling was 14.0 ($SD = 3.1$). Approximately 60% of the respon-

dents identified their job sector as wildlife/national park protection, 20% as nonprofit/NGO, and 20% as law enforcement, and overall they had an average of 9.6 years of experience working in wildlife protection ($SD = 6.1$).

Rangers from National Park, WWF, WCS, and YABI directly conduct field patrols over conservation areas, whereas the prosecutors report cases to lawyers and judges who

can open official investigations. Participants were native speakers of Bahasa Indonesia. The instructors delivered the workshop in English and interpreters translated all the material throughout the 3-day course. All written materials were made available in both English and Bahasa Indonesia.

Security Game Tutorials On the first day of the workshop, we introduced security game examples, beginning with a basic security game. We explained how the defender could optimally conduct patrols over targets and how attackers may respond against that strategy. We next covered (1) real-world security problems and applications of security games for protecting critical infrastructure and (2) challenges in wildlife protection and the application of security games to this domain, reviewing how adversaries' behaviors are modeled and how to optimize patrolling strategies through allocating limited security resources. Finally, we presented our PAWS software — which was built based on security game models for addressing wildlife protection problems — describing data inputs (e.g., animal density, poaching data) and what outputs are generated (e.g., models of poachers' behaviors and suggested patrolling routes for rangers). We described how similar approaches had been previously successfully used in the wildlife domain, and how they could be used on Sumatra via PAWS. This was our first introduction of PAWS to rangers in Indonesia, with the goal of obtaining feedback to improve PAWS and using it to generate patrols there in the future.

Discussion Sessions Participants engaged in several discussion sessions on topics including Resources (i.e., factors that motivate people to enter protected areas), Illegal Activities (i.e., types of illegal activities in conservation areas), and Wildlife Protection (i.e., improving security approaches). They exchanged knowledge about challenges in wildlife protection and generated potential solutions. For example, participants identified as a key challenge the existence of well-organized poacher groups who leverage information flows via local communities, obtaining information on ranger locations. We encouraged participants to develop solutions and provide feedback that could be quantified in game-theoretic models and incorporated into PAWS software. The ideas generated were valuable for specifying additional modeling components to incorporate in PAWS, including different groups of rangers/poachers and different



Figure 5: Workshop participants

types of animals, the role of local communities, and interactions between these components.

Role-Playing Games In addition to the tutorials and discussions, participants played the paired role-playing games as poachers and rangers (on computer and board games). Every participant played five rounds of the computer-based games as poachers. After each round, the poacher behavior models were updated based on participants' responses, and each subsequent game used a defender strategy created using these updated models. By playing these poaching games in a repeated fashion, the participants developed a better understanding of how poachers may react to rangers' strategies. They also learned about weaknesses of defender strategies, aiding them in better designing and testing patrol strategies in the future. Finally, participants learned how patrolling strategies generated by PAWS adapt and improve over time as more data are collected.

Additionally, we created physical board games which mimicked the computer-based game, replicating exactly the animal densities, rewards and penalties. In these board games, the participants were divided into two groups, each which in turn played as rangers (who created patrol strategies) and poachers (who decided where to poach in games generated by the other team). By playing these board games in both roles, participants learned how security games are created, how poachers might react to different defender strategies, and how AI software such as PAWS makes decisions based on models of players' behaviors.

On the final day, we presented the game results, i.e., the defender utilities based on poachers' decisions in the online games. The PAWS strategy improved over time, as indicated by increasing rangers' utility.

Results of Role Playing Games Fig. 6 shows the defender utilities obtained by deploying PAWS strategies over several rounds against workshop participants. In the figure, higher positive values along the y-axis indicate better defender results.

We observe that PAWS performs significantly better than a conservative game-theoretic strategy like Maximin, and completely outperforms the workshop participants.

Feedback All respondents rated the usefulness of the workshop as neutral or better, with 86% rating it as at least somewhat useful ($mean = 5.7$ on a scale of 1 [completely useless] – 7 [extremely useful], $SD = 0.8$). More than 96% of respondents reported that the workshop was at least somewhat important ($mean = 6.0$ on a scale of 1 [extremely unimportant] – 7 [extremely important], $SD = 1.0$). Additionally, more than 86% of respondents reported that they were at least somewhat likely to recommend the workshop to others in relevant fields ($mean = 6.04$ on a scale of 1




Figure 6: Defender utilities.

[strongly disagree] – 7 [strongly agree], $SD = 1.04$).

Based on open-ended responses, respondents largely supported the purpose of the PAWS software. Roughly half ($n = 15$) commented that it served to optimize patrols and would make the job of patrolling easier. When asked about PAWS software limitations, respondents recommended increasing complexity of models, including approximately one-third ($n = 10$) of respondents suggesting that dynamic animal distribution models be added. Respondents also suggested improving poaching predictions and that further research and testing be conducted prior to deployment.

Conclusion

This paper describes a novel approach that leveraged a game-based research platform to teach general AI and game-theoretic concepts to several audiences. These included students at an urban public high school, university undergraduate students, and law enforcement officers and park rangers who protect wildlife in Indonesia. The findings highlight the value of our game platform for teaching concepts in AI to learners representing a broad swath of age, gender, and cultural backgrounds.

Across these diverse audiences, a teaching method that proved to be effective was engaging learners in role-playing games. Members of all three audiences played the role of two different actors as part of our paired role-playing games, which included both computer and board games: (1) playing as rangers, they generated defender strategies to protect against poachers' attacks; (2) as poachers, they attempted to outsmart the defender strategies generated by their peers to earn the highest possible rewards. This methodology not only gave learners valuable hands-on experience in complex AI and game-theoretic concepts, but also in the development of real-world software applications for security. One limitation of our game is it is unclear how it would generalize to teaching AI topics beyond game theory and security games. In the future, we plan to develop activities focusing on other AI topics that can be integrated into our overall teaching approach. For instance, we will create an activity for learners to analyze defender strategies in depth, bringing to light subtle human biases that may impact initial strategies, thereby highlighting benefits of AI agents (compared to humans) in decision-making.

Participant feedback was consistently positive, with the majority of participants from all three audiences rating the learning experiences as useful, and in the case of the university and high school students, increasing their interest in AI. As well, most participants indicated they would recommend the learning activity to others. These results are particularly notable given the wide age range and cultural backgrounds of participants, and suggests that our approach was broadly accessible and engaging.

These findings also suggest that our approach could be applied successfully to additional audiences. A possible additional target group is security organizations who could benefit from security applications based on AI and game theory. Enhancing decision-makers' and field officers' understanding of the theory on which these applications are based could improve adoption rates of emerging AI-based decision aids.

Acknowledgements

This work was supported by MURI Grant W911NF-11-1-03 and World Wildlife Fund PAWS Anti-Poaching Grant.

References

- DeNero, J., and Klein, D. 2010. Teaching introductory artificial intelligence with pac-man. In *EAAI*.
- Fang, F.; Stone, P.; and Tambe, M. 2015. When security games go green: Designing defender strategies to prevent poaching and illegal fishing. In *IJCAI*.
- Haskell, W. B.; Kar, D.; Fang, F.; Tambe, M.; Cheung, S.; and Denicola, L. E. 2014. Robust protection of fisheries with COMPASS. In *IAAI*.
- Johnson, M. P.; Fang, F.; and Tambe, M. 2012. Patrol strategies to maximize pristine forest area. In *AAAI*.
- Kar, D.; Fang, F.; Fave, F. D.; Sintov, N.; and Tambe, M. 2015. a game of thrones: When human behavior models compete in repeated stackelberg security games. In *AAMAS*.
- Parsons, S., and Sklar, E. 2004. Teaching AI using LEGO mindstorms. In *AAAI Spring Symposium*.
- Pita, J.; Jain, M.; Western, C.; Portway, C.; Tambe, M.; Ordonez, F.; Kraus, S.; and Paruchuri, P. 2008. Deployed ARMOR protection: The application of a game theoretic model for security at the Los Angeles International Airport. In *AA-MAS*.
- 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 *AAMAS*.
- Stern, L., and Sterling, L. 1996. Teaching ai algorithms using animations reinforced by interactive exercises. In *Proceedings of the 2Nd Australasian Conference on Computer Science Education, ACSE '97*, 78–83. New York, NY, USA: ACM.
- Tambe, M.; Balsamo, A.; and Bowring, E. 2008. Using science fiction in teaching artificial intelligence. In *AAAI Spring Symposium*, 86–91.
- Tambe, M. 2011. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press.
- Wollowski, M. 2014. Teaching with watson. In *EAAI*, 3044–3049.
- Wong, D.; Zink, R.; and Koenig, S. 2010. Teaching artificial intelligence and robotics via games. In *EAAI*.
- Yang, R.; Ford, B.; Tambe, M.; and Lemieux, A. 2014. Adaptive resource allocation for wildlife protection against illegal poachers. In *AAMAS*.
- Yin, Z.; Korzhyk, D.; Kiekintveld, C.; Conitzer, V.; ; and Tambe, M. 2010. Stackelberg vs. nash in security games: Interchangeability, equivalence, and uniqueness. In *AAMAS*.