

Figure 5: Total number of crime




Figure 6: Individual Accuracy




Figure 7: Varying data




Figure 8: Varying data(Runtime)




Figure 9: Vary C




Figure 10: Compare with deployed




Figure 11: Vary T_u




Figure 12: Vary T_u (Runtime)

tic criminal’s behavior. However, the defender does not know the type of criminals in our experiments. Instead, the defender starts by executing a random patrol schedule for 300 steps and collects the corresponding crime report using which they learn an initial criminal behavior model. The criminal responds to the defenders’ patrol schedule as predicted by the behavior model in [20]. Since the criminal behavior in [20] is probabilistic, we run the experiment 30 times and each data point we report in this part is an average over these 30 instances. We fix the number of patrol officers to $2N - 2$, where N is the number of targets. This number is consistent with our real data-set numbers (8 officers for 5 targets), where there were enough officers to allocate one officer to each target, but not enough to allocate two officers to each target. We use EMC^2 algorithm as the learning algorithm.

Results: Figure 11 to 13 presents the results from our experiments about the online learning and planning mechanism. Four planning mechanisms that we consider are as follows: first, a random planning mechanism that randomly generates allocation strategy with limited resources; second, a pure planning mechanism, where we learn the criminal behavior model once and apply this model to plan for the entire horizon T using DOGS algorithm; third, a online planning mechanism with greedy planning algorithm that updates every T_u time-steps; and the last mechanism is online planning mechanism with DOGS algorithm that also updates every T_u time-steps. In Figure 11, the total planning horizon T is set to 600. In addition to the four planning mechanisms, we also consider the worst case where the defender always protect the least valuable targets. The x-axis shows the update interval T_u , which is the time interval after which we update criminals’ behavior model. The y-axis is the expected number of crimes that happens under the deployed allocation strategy within 600 steps. Expected number of crimes under pure planning mechanism stay the same with different T_u because it does not update the criminals’ model at all. For online mechanisms, the expected number of crimes increases




Figure 13: Varying N

as the update interval T_u increases. This is because with infrequent updates of the criminals’ behavior model, we cannot keep up with the real criminals’ behavior. In addition, with any size of the update interval, DOGS algorithm outperforms the greedy algorithm. In Figure 12, we present the runtime of three mechanisms for the same experiment. We do not show the runtime for the random planning mechanism as it is small and same for any planning horizon T . The runtime decreases as the update interval T_u increases. There is a runtime-quality trade-off in choosing T_u . Figure 13 shows the performance of the four planning mechanisms, but with different number of targets in the model. The x-axis is the number of targets in the graph and the y-axis is the expected number of crimes under the deployed strategy. We set $T = 600$, $T_u = 2$. The results here are similar to the results of Fig. 11.

These results lead us to conclude that online mechanisms outperform the baseline planning mechanisms significantly in any settings. For online mechanisms, DOGS achieves better performance while greedy planning algorithm requires less runtime. Thus, based on the specific problem being solved, the appropriate algorithm must be chosen judiciously.

These results lead us to conclude that online mechanisms outperform the baseline planning mechanisms significantly in any settings. For online mechanisms, DOGS achieves better performance while greedy planning algorithm requires less runtime. Thus, based on the specific problem being solved, the appropriate algorithm must be chosen judiciously.

8. CONCLUSION

This paper introduces a novel framework to design patrol allocation against adaptive opportunistic criminals. First, we model the interaction between officers and adaptive opportunistic criminals as a DBN. Next, we propose a sequence of modifications to the basic DBN resulting in a compact model that enables better learning accuracy and running time. Finally, we present an iterative learning and planning mechanism with two planning algorithm to keep pace with adaptive opportunistic criminals. Experimental validation with real data supports our choice of model and assumptions. Further, our modeling assumptions were informed by inputs from our collaborators in the DPS at USC. These promising results have opened up the possibility of deploying our method in USC. This paper has further opened up the integration of opportunistic crime security games [20] with machine learning.

9. ACKNOWLEDGEMENT

This research is supported by MURI grant W911NF-11-1-0332. We thank Kevin Knight for many helpful comments on this work.

REFERENCES

- [1] S. M. Aji and R. J. McEliece. The generalized distributive law. *Information Theory, IEEE Transactions on*, 46(2):325–343, 2000.
- [2] N. Basilico, N. Gatti, and F. Amigoni. Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pages 57–64. International Foundation for Autonomous Agents and Multiagent Systems, 2009.
- [3] N. Basilico, N. Gatti, T. Rossi, S. Ceppi, and F. Amigoni. Extending algorithms for mobile robot patrolling in the presence of adversaries to more realistic settings. In *Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology-Volume 02*, pages 557–564. IEEE Computer Society, 2009.
- [4] C. M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006.
- [5] A. Blum, N. Haghtalab, and A. D. Procaccia. Learning optimal commitment to overcome insecurity. In *Proceedings of the 28th Annual Conference on Neural Information Processing Systems (NIPS). Forthcoming*, 2014.
- [6] A. Blumer, A. Ehrenfeucht, D. Haussler, and M. K. Warmuth. Occam’s razor. *Inf. Process. Lett.*, 24(6):377–380, Apr. 1987.
- [7] X. Boyen and D. Koller. Tractable inference for complex stochastic processes. In *Proceedings of the Fourteenth conference on Uncertainty in artificial intelligence*, pages 33–42. Morgan Kaufmann Publishers Inc., 1998.
- [8] H. Chen, W. Chung, J. J. Xu, G. Wang, Y. Qin, and M. Chau. Crime data mining: a general framework and some examples. *Computer*, 37(4):50–56, 2004.
- [9] J. S. De Bruin, T. K. Cocx, W. A. Kusters, J. F. Laros, and J. N. Kok. Data mining approaches to criminal career analysis. In *Data Mining, 2006. ICDM’06. Sixth International Conference on*, pages 171–177. IEEE, 2006.
- [10] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 1–38, 1977.
- [11] J. P. Hespanha, M. Prandini, and S. Sastry. Probabilistic pursuit-evasion games: A one-step nash approach. In *Decision and Control, 2000. Proceedings of the 39th IEEE Conference on*, volume 3, pages 2272–2277. IEEE, 2000.
- [12] M. Jain, J. Tsai, J. Pita, C. Kiekintveld, S. Rathi, M. Tambe, and F. Ordóñez. Software assistants for randomized patrol planning for the lax airport police and the federal air marshal service. *Interfaces*, 40(4):267–290, 2010.
- [13] A. X. Jiang, Z. Yin, C. Zhang, M. Tambe, and S. Kraus. Game-theoretic randomization for security patrolling with dynamic execution uncertainty. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*, pages 207–214. International Foundation for Autonomous Agents and Multiagent Systems, 2013.
- [14] S. V. Nath. Crime pattern detection using data mining. In *Web Intelligence and Intelligent Agent Technology Workshops, 2006. WI-IAT 2006 Workshops. 2006 IEEE/WIC/ACM International Conference on*, pages 41–44. IEEE, 2006.
- [15] G. Oatley, B. Ewart, and J. Zeleznikow. Decision support systems for police: Lessons from the application of data mining techniques to soft forensic evidence. *Artificial Intelligence and Law*, 14(1-2):35–100, 2006.
- [16] E. Shieh, B. An, R. Yang, M. Tambe, C. Baldwin, J. DiRenzo, B. Maule, and G. Meyer. Protect: A deployed game theoretic system to protect the ports of the united states. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pages 13–20. International Foundation for Autonomous Agents and Multiagent Systems, 2012.
- [17] M. B. Short, M. R. D’ORSOGNA, V. B. Pasour, G. E. Tita, P. J. Brantingham, A. L. Bertozzi, and L. B. Chayes. A statistical model of criminal behavior. *Mathematical Models and Methods in Applied Sciences*, 18(supp01):1249–1267, 2008.
- [18] M. Tambe. *Security and game theory: algorithms, deployed systems, lessons learned*. Cambridge University Press, 2011.
- [19] R. Yang, B. Ford, M. Tambe, and A. Lemieux. Adaptive resource allocation for wildlife protection against illegal poachers. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 453–460. International Foundation for Autonomous Agents and Multiagent Systems, 2014.
- [20] C. Zhang, A. X. Jiang, M. B. Short, P. J. Brantingham, and M. Tambe. Defending against opportunistic criminals: New game-theoretic frameworks and algorithms. In *Decision and Game Theory for Security*, pages 3–22. Springer, 2014.