



















## REFERENCES

- [1] N. Akbarzadeh and A. Mahajan. 2019. Restless bandits with controlled restarts: Indexability and computation of Whittle index. In *2019 IEEE Conference on Decision and Control*. IEEE.
- [2] Biswarup Bhattacharya. 2018. Restless bandits visiting villages: A preliminary study on distributing public health services. In *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*. 1–8.
- [3] Arpita Biswas, Gaurav Aggarwal, Pradeep Varakantham, and Milind Tambe. 2021. Learn to Intervene: An Adaptive Learning Policy for Restless Bandits in Application to Preventive Healthcare. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence*.
- [4] Arpita Biswas, Shweta Jain, Debmalya Mandal, and Y Narahari. 2015. A truthful budget feasible multi-armed bandit mechanism for crowdsourcing time critical tasks. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*. 1101–1109.
- [5] J Nell Brownstein, Farah M Chowdhury, Susan L Norris, Tanya Horsley, Leonard Jack Jr, Xuanping Zhang, and Dawn Satterfield. 2007. Effectiveness of community health workers in the care of people with hypertension. *American journal of preventive medicine* 32, 5 (2007), 435–447.
- [6] Alicia H Chang, Andrea Polesky, and Gulshan Bhatia. 2013. House calls by community health workers and public health nurses to improve adherence to isoniazid monotherapy for latent tuberculosis infection: a retrospective study. *BMC public health* 13, 1 (2013), 894.
- [7] K.D. Glazebrook, D. Ruiz-Hernandez, and C. Kirkbride. 2006. Some indexable families of restless bandit problems. *Adv. Appl. Probab* (2006), 643–672.
- [8] Jeffrey Thomas Hawkins. 2003. *A Lagrangian decomposition approach to weakly coupled dynamic optimization problems and its applications*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- [9] Y. Hsu. 2018. Age of information: Whittle index for scheduling stochastic arrivals. In *2018 IEEE International Symposium on Information Theory*. IEEE.
- [10] Weici Hu and Peter Frazier. 2017. An asymptotically optimal index policy for finite-horizon restless bandits. *arXiv preprint arXiv:1707.00205* (2017).
- [11] S. Jiang, Z. Song, O. Weinstein, and H. Zhang. 2020. Faster dynamic matrix inverse for faster lps. *arXiv preprint arXiv:2004.07470* (2020).
- [12] Varun Kanade, H Brendan McMahan, and Brent Bryan. 2009. Sleeping experts and bandits with stochastic action availability and adversarial rewards. In *Artificial Intelligence and Statistics*. PMLR, 272–279.
- [13] J. A. Killian, B. Wilder, A. Sharma, V. Choudhary, B. Dilkina, and M. Tambe. 2019. Learning to Prescribe Interventions for Tuberculosis Patients using Digital Adherence Data. In *KDD*.
- [14] Robert Kleinberg, Alexandru Niculescu-Mizil, and Yogeshwer Sharma. 2010. Regret bounds for sleeping experts and bandits. *Machine learning* 80, 2 (2010), 245–272.
- [15] Jerome Le Ny, Munther Dahleh, and Eric Feron. 2008. Multi-UAV dynamic routing with partial observations using restless bandit allocation indices. In *2008 American Control Conference*. IEEE, 4220–4225.
- [16] Elliot Lee, Mariel S Lavieri, and Michael Volk. 2019. Optimal screening for hepatocellular carcinoma: A restless bandit model. *Manufacturing & Service Operations Management* 21, 1 (2019), 198–212.
- [17] John DC Little. 1961. A proof for the queuing formula:  $L = \lambda W$ . *Operations research* 9, 3 (1961), 383–387.
- [18] K. Liu and Q. Zhao. 2010. Indexability of restless bandit problems and optimality of Whittle index for dynamic multichannel access. *IEEE Transactions on Information Theory* 56, 11 (2010), 5547–5567.
- [19] K. Liu and Q. Zhao. 2010. Indexability of restless bandit problems and optimality of Whittle index for dynamic multichannel access. *IEEE Transactions on Information Theory* (2010), 5547–5567.
- [20] Bernd Löwe, Jürgen Unützer, Christopher M Callahan, Anthony J Perkins, and Kurt Kroenke. 2004. Monitoring depression treatment outcomes with the patient health questionnaire-9. *Medical care* (2004), 1194–1201.
- [21] Aditya Mate, Jackson A Killian, Haifeng Xu, Andrew Perrault, and Milind Tambe. 2020. Collapsing Bandits and Their Application to Public Health Interventions. In *Advances in Neural and Information Processing Systems (NeurIPS)*.
- [22] Aditya Mate, Lovish Madaan, Aparna Taneja, Neha Madhiwalla, Shresth Verma, Gargi Singh, Aparna Hegde, Pradeep Varakantham, and Milind Tambe. 2021. Field Study in Deploying Restless Multi-Armed Bandits: Assisting Non-Profits in Improving Maternal and Child Health. *arXiv preprint arXiv:2109.08075* (2021).
- [23] Aditya Mate, Andrew Perrault, and Milind Tambe. 2021. Risk-Aware Interventions in Public Health: Planning with Restless Multi-Armed Bandits. In *International Conference on Autonomous Agents and Multiagent Systems*.
- [24] Nicolas Meuleau, Milos Hauskrecht, Kee-Eung Kim, Leonid Peshkin, Leslie Pack Kaelbling, Thomas L Dean, and Craig Boutilier. 1998. Solving very large weakly coupled Markov decision processes. In *AAAI/IAAI*. 165–172.
- [25] Christopher Mundorf, Arti Shankar, Tracy Moran, Sherry Heller, Anna Hassan, Emily Harville, and Maureen Lichtveld. 2018. Reducing the risk of postpartum depression in a low-income community through a community health worker intervention. *Maternal and child health journal* 22, 4 (2018), 520–528.
- [26] Patrick M Newman, Molly F Franke, Jafet Arrieta, Hector Carrasco, Patrick Elliott, Hugo Flores, Alexandra Friedman, Sophia Graham, Luis Martinez, Lindsay Palazuelos, et al. 2018. Community health workers improve disease control and medication adherence among patients with diabetes and/or hypertension in Chiapas, Mexico: an observational stepped-wedge study. *BMJ Global Health* (2018).
- [27] José Nino-Mora. 2011. Computing a classic index for finite-horizon bandits. *INFORMS Journal on Computing* 23, 2 (2011), 254–267.
- [28] Christos H Papadimitriou and John N Tsitsiklis. 1994. The complexity of optimal queueing network control. In *Proceedings of IEEE 9th Annual Conference on Structure in Complexity Theory*. IEEE, 318–322.
- [29] Y. Qian, C. Zhang, B. Krishnamachari, and M. Tambe. 2016. Restless poachers: Handling exploration-exploitation tradeoffs in security domains. In *AAMAS*.
- [30] Jane Rahedi Ong’ang’o, Christina Mwachari, Hillary Kipruto, and Simon Karanja. 2014. The effects on tuberculosis treatment adherence from utilising community health workers: a comparison of selected rural and urban settings in Kenya. *PLoS One* 9, 2 (2014), e88937.
- [31] B. Sombabu, A. Mate, D. Manjunath, and S. Moharir. 2020. Whittle index for AoI-aware scheduling. In *IEEE International Conference on Communication Systems & Networks (COMSNETS)*. IEEE.
- [32] P. Whittle. 1988. Restless bandits: Activity allocation in a changing world. *J. Appl. Probab.* 25, A (1988), 287–298.
- [33] Gabriel Zayas-Caban, Stefanus Jasin, and Guihua Wang. 2019. An asymptotically optimal heuristic for general nonstationary finite-horizon restless multi-armed, multi-action bandits. *Advances in Applied Probability* 51, 3 (2019), 745–772.