

REFERENCES

- [1] Kai Arulkumaran, Marc Peter Deisenroth, Miles Brundage, and Anil Anthony Bharath. 2017. Deep reinforcement learning: A brief survey. *IEEE Signal Processing Magazine* 34, 6 (2017), 26–38.
- [2] F. G. Ball, E. S. Knock, and P. D. O’Neill. 2015. Stochastic epidemic models featuring contact tracing with delays. *Math biosciences* 266 (2015), 23–35.
- [3] A. Banerjee, A. G. Chandrasekhar, E. Duflo, and M. O. Jackson. 2013. The diffusion of microfinance. *Science* 341, 6144 (2013), 1236498.
- [4] Irwan Bello, Hieu Pham, Quoc V Le, Mohammad Norouzi, and Samy Bengio. 2016. Neural combinatorial optimization with reinforcement learning. *arXiv preprint arXiv:1611.09940* (2016).
- [5] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *ICML*. 41–48.
- [6] Marian Boguná, Claudio Castellano, and Romualdo Pastor-Satorras. 2013. Nature of the epidemic threshold for the susceptible-infected-susceptible dynamics in networks. *Physical review letters* 111, 6 (2013), 068701.
- [7] Hanjun Dai, Bo Dai, and Le Song. 2016. Discriminative embeddings of latent variable models for structured data. In *ICML*. 2702–2711.
- [8] Peter Dayan and Geoffrey E Hinton. 1993. Feudal reinforcement learning. In *NeurIPS*. 271–278.
- [9] Thomas G Dietterich. 2000. Hierarchical reinforcement learning with the MAXQ value function decomposition. *Journal of AI research* 13 (2000), 227–303.
- [10] Kimon Drakopoulos, Asuman Ozdaglar, and John N Tsitsiklis. 2014. An efficient curing policy for epidemics on graphs. *IEEE TNSE* 1, 2 (2014), 67–75.
- [11] Kimon Drakopoulos, Asuman Ozdaglar, and John N Tsitsiklis. 2016. When is a network epidemic hard to eliminate? *Mathematics of Operations Research* (2016).
- [12] Thomas DuBois, Stephen Eubank, and Aravind Srinivasan. 2012. The effect of random edge removal on network degree sequence. *Electronic journal of combinatorics* 19, 1 (2012).
- [13] Ken TD Eames and Matt J Keeling. 2003. Contact tracing and disease control. *Proceedings of the Royal Society of London. Series B: Biological Sciences* 270, 1533 (2003), 2565–2571.
- [14] Luca Ferretti, Chris Wymant, Michelle Kendall, Lele Zhao, Anel Nurtay, Lucia Abeler-Dörner, Michael Parker, David Bonsall, and Christophe Fraser. 2020. Quantifying SARS-CoV-2 transmission suggests epidemic control with digital contact tracing. *Science* 368, 6491 (2020).
- [15] Johan Giesecke. 2017. *Modern infectious disease epidemiology*. CRC Press.
- [16] Alexander K Hartmann and Martin Weigt. 2005. *Phase transitions in combinatorial optimization problems*. Vol. 67. Wiley Online Library.
- [17] Jessica Hoffmann and Constantine Caramanis. 2018. The Cost of Uncertainty in Curing Epidemics. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 2, 2 (2018), 31.
- [18] Petter Holme. 2004. Efficient local strategies for vaccination and network attack. *EPL (Europhysics Letters)* 68, 6 (2004), 908.
- [19] Alex Irpan. 2018. Deep Reinforcement Learning Doesn’t Work Yet. <https://www.alexirpan.com/2018/02/14/r1-hard.html>.
- [20] L. Isella, J. Stehlé, A. Barrat, C. Cattuto, J.-F. Pinton, and W. Van den Broeck. 2011. What’s in a crowd? Analysis of face-to-face behavioral networks. *J. Theor. Biol.* (2011), 166.
- [21] David S Johnson. 1974. Approximation algorithms for combinatorial problems. *Journal of computer and system sciences* 9, 3 (1974), 256–278.
- [22] Leslie Pack Kaelbling, Michael L Littman, and Andrew W Moore. 1996. Reinforcement learning: A survey. *Journal of artificial intelligence research* 4 (1996), 237–285.
- [23] Harshavardhan Kamarthi, Priyesh Vijayan, Bryan Wilder, Balaraman Ravindran, and Milind Tambe. 2019. Learning policies for Social network discovery with Reinforcement learning. *arXiv preprint arXiv:1907.11625* (2019).
- [24] Harshavardhan Kamarthi, Priyesh Vijayan, Bryan Wilder, Balaraman Ravindran, and Milind Tambe. 2020. Influence Maximization in Unknown Social Networks: Learning Policies for Effective Graph Sampling. In *AAMAS*. 575–583.
- [25] David Kempe, Jon Kleinberg, and Éva Tardos. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the 9th ACM SIGKDD*. 137–146.
- [26] Elias Khalil, Hanjun Dai, Yuyu Zhang, Bistra Dilkina, and Le Song. 2017. Learning combinatorial optimization algorithms over graphs. In *NeurIPS*. 6348–6358.
- [27] Thomas N Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *ICLR-17*, Toulon.
- [28] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *NeurIPS*. 1097–1105.
- [29] Cong Li, Huijuan Wang, and Piet Van Mieghem. 2012. Degree and principal eigenvectors in complex networks. In *ICRN*. Springer, 149–160.
- [30] Fredrik Liljeros, Christofer R Edling, and Luis A Nunes Amaral. 2003. Sexual networks: implications for the transmission of sexually transmitted infections. *Microbes and infection* 5, 2 (2003), 189–196.
- [31] Dermot Maher, Chris Dye, Katherine Floyd, Andrea Pantoja, Knut Lonnroth, Alasdair Reid, Eva Nathanson, Thad Pennas, Uli Fruth, Jane Cunningham, et al. 2007. Planning to improve global health: the next decade of tuberculosis control. *Bulletin of the World Health Organization* 85 (2007), 341–347.
- [32] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602* (2013).
- [33] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *nature* 518, 7540 (2015), 529–533.
- [34] Han-Ching Ou, Arunesh Sinha, Sze-Chuan Suen, Andrew Perrault, Alpan Raval, and Milind Tambe. 2020. Who and When to Screen: Multi-Round Active Screening for Network Recurrent Infectious Diseases Under Uncertainty. In *AAMAS*. 992–1000.
- [35] Han Ching Ou, Kai Wang, Finale Doshi-Velez, and Milind Tambe. [n.d.]. Active Screening on Recurrent Diseases Contact Networks with Uncertainty: a Reinforcement Learning Approach. ([n.d.]).
- [36] P. Panzarasa, T. Opsahl, and K. M. Carley. 2009. Patterns and dynamics of users’ behavior and interaction: Network analysis of an online community. *J. Assoc. Inf. Sci. Technol.* (2009), 911.
- [37] Ronald Parr and Stuart J Russell. 1998. Reinforcement learning with hierarchies of machines. In *NeurIPS*. 1043–1049.
- [38] Roni Parshani, Shai Carmi, and Shlomo Havlin. 2010. Epidemic threshold for the susceptible-infectious-susceptible model on random networks. *Physical review letters* 104, 25 (2010), 258701.
- [39] Romualdo Pastor-Satorras and Alessandro Vespignani. 2001. Epidemic spreading in scale-free networks. *Physical review letters* 86, 14 (2001), 3200.
- [40] Wei Qiu, Haipeng Chen, and Bo An. 2019. Dynamic Electronic Toll Collection via Multi-Agent Deep Reinforcement Learning with Edge-Based Graph Convolutional Networks. In *IJCAI*. 4568–4574.
- [41] Yizhi Ren, Mengjin Jiang, Ye Yao, Ting Wu, Zhen Wang, Mengkun Li, and Kim-Kwang Raymond Choo. 2018. Node Immunization in Networks with Uncertainty. In *(TrustCom/BigDataSE)*. IEEE, 1392–1397.
- [42] S. Saha, A. Adiga, B. A. Prakash, and A. K. S. Vullikanti. 2015. Approximation algorithms for reducing the spectral radius to control epidemic spread. In *Proceedings of the 2015 SIAM ICDM*. SIAM, 568–576.
- [43] M. Salathé, M. Kazandjieva, J. W. Lee, P. Levis, M. W. Feldman, and J. H. Jones. 2010. A high-resolution human contact network for infectious disease transmission. *Proceedings of the National Academy of Sciences* 107, 51 (2010), 22020–22025.
- [44] K. Scaman, A. Kalogeratos, and N. Vayatis. 2016. Suppressing epidemics in networks using priority planning. *IEEE TNSE* (2016).
- [45] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. 2016. Mastering the game of Go with deep neural networks and tree search. *nature* 529, 7587 (2016), 484–489.
- [46] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. 2017. Mastering the game of go without human knowledge. *nature* 550, 7676 (2017), 354–359.
- [47] Jeffrey A Smith and James Moody. 2013. Structural effects of network sampling coverage I: Nodes missing at random. *Social networks* 35, 4 (2013), 652–668.
- [48] Hyungseok Song, Hyeryung Jang, Hai H Tran, Se-eun Yoon, Kyunghwan Son, Donggyu Yun, Hyouju Chung, and Yung Yi. 2019. Solving continual combinatorial selection via deep reinforcement learning. *arXiv preprint arXiv:1909.03638* (2019).
- [49] C. Sun and Y.-H. Hsieh. 2010. Global analysis of an SEIR model with varying population size and vaccination. *Applied Mathematical Modelling* 34, 10 (2010), 2685–2697.
- [50] Richard S Sutton et al. 1998. *Introduction to reinforcement learning*. Vol. 135. MIT press Cambridge.
- [51] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- [52] Richard S Sutton, Doina Precup, and Satinder Singh. 1999. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence* 112, 1-2 (1999), 181–211.
- [53] David Taylor-Robinson. 1994. Chlamydia trachomatis and sexually transmitted disease.
- [54] H. Tong, B. A. Prakash, T. Eliassi-Rad, M. Faloutsos, and C. Faloutsos. 2012. Gelling, and melting, large graphs by edge manipulation. In *Proceedings of the 21st ACM international CIKM*. ACM, 245–254.
- [55] P. Vanhems et al. 2013. Estimating potential infection transmission routes in hospital wards using wearable proximity sensors. *PLoS one* 8 (2013), 73970.
- [56] Y. Wang, C. Chakrabarti, D. Wang, and C. Faloutsos. 2003. Epidemic spreading in real networks: An eigenvalue viewpoint. In *SRDS*. IEEE, 25–34.
- [57] Christopher JCH Watkins and Peter Dayan. 1992. Q-learning. *Machine learning* 8, 3-4 (1992), 279–292.
- [58] Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. 2014. Recurrent neural network regularization. *arXiv preprint arXiv:1409.2329* (2014).
- [59] Yao Zhang and B Aditya Prakash. 2015. Data-aware vaccine allocation over large networks. *ACM TKDD* 10, 2 (2015), 1–32.