

## Textbooks

1. Reinforcement Learning: An Introduction, by Richard S. Sutton and Andrew G. Barto
2. Speech and Language Processing (3rd edition), by Daniel Jurafsky and James H. Martin

## Papers

1. Ecoffet, A., Huizinga, J., Lehman, J. et al. First return, then explore. *Nature* 590, 580–586 (2021). (<https://www.nature.com/articles/s41586-020-03157-9>)
2. Agarwal, A., Henaff, M., Kakade, S., & Sun, W. (2020). Pc-pg: Policy cover directed exploration for provable policy gradient learning. *Advances in Neural Information Processing Systems*, 33, 13399-13412. (<https://arxiv.org/abs/2007.08459>)
3. Kearns, M., Singh, S. Near-Optimal Reinforcement Learning in Polynomial Time. *Machine Learning* 49, 209–232 (2002). (<https://www.cis.upenn.edu/~mkearns/papers/KearnsSinghE3.pdf>)
4. Narasimhan, K., Kulkarni, T., & Barzilay R. 2015. Language Understanding for Text-based Games using Deep Reinforcement Learning. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1–11, Lisbon, Portugal. Association for Computational Linguistics. (<https://aclanthology.org/D15-1001/>)
5. He, J., Chen, J., He, X., Gao, J., Li, L., Deng, L., & Ostendorf, M. 2016. Deep Reinforcement Learning with a Natural Language Action Space. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1621–1630, Berlin, Germany. Association for Computational Linguistics. (<https://arxiv.org/abs/1511.04636>)
6. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30. (<https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>)
7. Devlin, J., Chang, M., Lee, K., and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
8. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. *Advances in neural information processing systems*, 33, 1877-1901.
9. Yao, S., Narasimhan, K., & Hausknecht, M. 2021. Reading and Acting while Blindfolded: The Need for Semantics in Text Game Agents. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational*

*Linguistics: Human Language Technologies*, pages 3097–3102, Online. Association for Computational Linguistics. (<https://aclanthology.org/2021.naacl-main.247/>)

10. Hausknecht, M., Ammanabrolu, P., Côté, M. A., & Yuan, X. (2020, April). Interactive fiction games: A colossal adventure. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 34, No. 05, pp. 7903-7910). (<https://ojs.aaai.org/index.php/AAAI/article/view/6297>)
11. Pathak, D., Agrawal, P., Efros, A. A., & Darrell, T. (2017, July). Curiosity-driven exploration by self-supervised prediction. In *International conference on machine learning* (pp. 2778-2787). PMLR. (<https://proceedings.mlr.press/v70/pathak17a.html>)
12. Luketina, J., Nardelli, N., Farquhar, G., Foerster, J.N., Andreas, J., Grefenstette, E., Whiteson, S., & Rocktäschel, T. (2019). A Survey of Reinforcement Learning Informed by Natural Language. *IJCAI*. (<https://arxiv.org/pdf/1906.03926.pdf>)