

REFERENCES

- [1] Marc G. Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. 2013. The Arcade Learning Environment: An Evaluation Platform for General Agents. *Journal of Artificial Intelligence Research* 47 (Jun 2013), 253–279.
- [2] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. OpenAI Gym. (2016). arXiv:cs.LG/1606.01540
- [3] Antonion R. Damasio. 1989. Time-locked multiregional retroactivation: A systems-level proposal for the neural substrates of recall and recognition. *Cognition* 33, 1-2 (1989), 25–62.
- [4] Chelsea Finn, Xin Yu Tan, Yan Duan, Trevor Darrell, Sergey Levine, and Pieter Abbeel. 2016. Deep spatial autoencoders for visuomotor learning. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 512–519.
- [5] Shani Gamrian and Yoav Goldberg. 2019. Transfer Learning for Related Reinforcement Learning Tasks via Image-to-Image Translation. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, Vol. 97. PMLR, 2063–2072.
- [6] Carles Gelada, Saurabh Kumar, Jacob Buckman, Ofir Nachum, and Marc G. Bellemare. 2019. DeepMDP: Learning Continuous Latent Space Models for Representation Learning. (2019). arXiv:cs.LG/1906.02736
- [7] Irina Higgins, Arka Pal, Andrei Rusu, Loic Matthey, Christopher Burgess, Alexander Pritzel, Matthew Botvinick, Charles Blundell, and Alexander Lerchner. 2017. DARLA: Improving Zero-Shot Transfer in Reinforcement Learning. In *Proceedings of the 34th International Conference on Machine Learning*, Vol. 70. PMLR, 1480–1490.
- [8] Diederik P. Kingma and Max Welling. 2013. Auto-Encoding Variational Bayes. (2013). arXiv:stat.ML/1312.6114
- [9] Timo Korthals. 2019. M^2 VAE - Derivation of a Multi-Modal Variational Autoencoder Objective from the Marginal Joint Log-Likelihood. (2019). arXiv:cs.LG/1903.07303
- [10] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. 2015. Continuous control with deep reinforcement learning. (2015). arXiv:cs.LG/1509.02971
- [11] Kaspar Meyer and Antonio Damasio. 2009. Convergence and divergence in a neural architecture for recognition and memory. *Trends in Neurosciences* 32, 7 (2009), 376–382.
- [12] 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.
- [13] Prajit Ramachandran, Barret Zoph, and Quoc V. Le. 2017. Searching for Activation Functions. (2017). arXiv:cs.NE/1710.05941
- [14] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. 2014. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. (2014). arXiv:stat.ML/1401.4082
- [15] Richard Sutton and Andrew Barto. 1998. *Reinforcement Learning: An Introduction*. MIT press Cambridge.
- [16] Masahiro Suzuki, Kotaro Nakayama, and Yutaka Matsuo. 2016. Joint Multimodal Learning with Deep Generative Models. (2016). arXiv:stat.ML/1611.01891
- [17] Christopher Watkins. 1989. *Learning from delayed rewards*. Ph.D. Dissertation. Cambridge University.
- [18] Hang Yin, Francisco S. Melo, Aude Billard, and Ana Paiva. 2017. Associate Latent Encodings in Learning from Demonstrations. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*. AAAI Press, 3848–3854.
- [19] Marvin Zhang, Sharad Vikram, Laura Smith, Pieter Abbeel, Matthew J. Johnson, and Sergey Levine. 2018. SOLAR: Deep Structured Representations for Model-Based Reinforcement Learning. (2018). arXiv:cs.LG/1808.09105