





## REFERENCES

- [1] Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. 2007. Analysis of representations for domain adaptation. *Advances in Neural Information Processing Systems (NIPS)* 19 (2007), 137–144.
- [2] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. *Proc. ICML* (2009), 41–48. <https://doi.org/10.1145/1553374.1553380>
- [3] Edwin V Bonilla, Felix V Agakov, and Christopher Williams. 2007. Kernel multi-task learning using task-specific features. *International Conference on Artificial Intelligence and Statistics* (2007), 43–50.
- [4] Rich Caruana. 1997. Multitask Learning. *Machine Learning* 28 (1997), 41–75.
- [5] Bruno Da Silva, George Konidaris, and Andrew Barto. 2012. Learning parameterized skills. *arXiv:1206.6398* (2012).
- [6] Eric Eaton, Marie desJardins, and Terran Lane. 2008. Modeling transfer relationships between learning tasks for improved inductive transfer. In *Joint European Conf. on Mach. Learn. and Knowledge Discovery in Databases*. Springer, 317–332.
- [7] Nick Hawes. 2011. A survey of motivation frameworks for intelligent systems. *Art. Intel. J.* 175, 5-6 (April 2011), 1020–1036.
- [8] David Isele and Akansel Cosgun. 2018. Selective Experience Replay for Lifelong Learning. *Proc. AAAI* (2018).
- [9] David Isele, Mohammad Rostami, and Eric Eaton. 2016. Using task features for zero-shot knowledge transfer in lifelong learning. *Proc. IJCAI* (2016).
- [10] Matthew Johnson, Katja Hofmann, Tim Hutton, and David Bignell. 2016. The Malmo Platform for Artificial Intelligence Experimentation.. In *Proc. IJCAI*. 4246–4247.
- [11] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy P Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. 2016. Asynchronous methods for deep reinforcement learning. In *Proc. ICML*.
- [12] Sanmit Narvekar, Jivko Sinapov, Matteo Leonetti, and Peter Stone. 2016. Source task creation for curriculum learning. In *Proc. AAMAS*. 566–574.
- [13] Sanmit Narvekar, Jivko Sinapov, and Peter Stone. 2017. Autonomous Task Sequencing for Customized Curriculum Design in Reinforcement Learning. (2017).
- [14] Doina Precup. 2000. Temporal abstraction in reinforcement learning. (2000).
- [15] Mark B Ring. 1997. CHILd: A first step towards continual learning. *Machine Learning* 28, 1 (1997), 77–104.
- [16] Jivko Sinapov, Sanmit Narvekar, Matteo Leonetti, and Peter Stone. 2015. Learning inter-task transferability in the absence of target task samples. In *Proc. AAMAS*. 725–733.
- [17] Aravind Srinivas, Sahil Sharma, and Balaraman Ravindran. 2017. Dynamic Action Repetition For Deep Reinforcement Learning. *Proc. AAAI* (2017).
- [18] 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.
- [19] Maxwell Svetlik, Matteo Leonetti, Jivko Sinapov, Rishi Shah, Nick Walker, and Peter Stone. 2017. Automatic Curriculum Graph Generation for Reinforcement Learning Agents.. In *Proc. AAAI*. 2590–2596.
- [20] S. Vattam, M. Klenk, M. Molineaux, and D.W. Aha. 2013. Breadth of approaches to goal reasoning: A research survey. In *Goal Reasoning: Papers from the ACS Workshop (Technical Report CS-TR-5029)*. College Park, MD: University of Maryland, Department of Computer Science., 222–231.
- [21] Joannes Vermorel and Mehryar Mohri. 2005. Multi-armed bandit algorithms and empirical evaluation. In *ECML*, Vol. 3720. Springer, 437–448.