

REFERENCES

- [1] M. E. Akintunde, A. Kevorchian, A. Lomuscio, and E. Pirovano. 2019. Verification of RNN-Based Neural Agent-Environment Systems. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI19)*. AAAI Press, 6006–6013.
- [2] M. E. Akintunde, A. Lomuscio, L. Maganti, and E. Pirovano. 2018. Reachability Analysis for Neural Agent-Environment Systems. In *Proceedings of the 16th International Conference on Principles of Knowledge Representation and Reasoning (KR18)*. AAAI Press, 184–199.
- [3] A. Biere, A. Cimatti, E. Clarke, O. Strichman, and Y. Zhu. 2003. Bounded Model Checking. *Advances in Computers* 58 (2003), 117–148.
- [4] R. H. Bordini, M. Fisher, W. Visser, and M. Wooldridge. 2006. Verifying Multi-agent Programs by Model Checking. *Autonomous Agents and Multi-Agent Systems* 12, 2 (2006), 239–256.
- [5] R. R. Bunel, I. Turkaslan, P. Torr, P. Kohli, and P. K. Mudigonda. 2018. A Unified View of Piecewise Linear Neural Network Verification. In *Proceedings of the 31st Annual Conference on Neural Information Processing Systems (NeurIPS18)*. Curran Associates, Inc., 4790–4799.
- [6] E. Clarke, A. Biere, R. Raimi, and Y. Zhu. 2001. Bounded Model Checking Using Satisfiability Solving. *Formal Methods in System Design* 19, 1 (2001), 7–34.
- [7] E. M. Clarke, O. Grumberg, and D. A. Peled. 1999. *Model Checking*. The MIT Press, Cambridge, Massachusetts.
- [8] C. D’Ambrosio, A. Lodi, and S. Martello. 2010. Piecewise linear approximation of functions of two variables in MLP models. *Operations Research Letters* 38, 1 (2010), 39–46.
- [9] T. T. Doan, Y. Yao, N. Alechina, and B. Logan. 2014. Verifying heterogeneous multi-agent programs. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS14)*. IFAAMAS, 149–156.
- [10] S. Dutta, X. Chen, and S. Sankaranarayanan. 2019. Reachability Analysis for Neural Feedback Systems using Regressive Polynomial Rule Inference. In *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control (HSCC19)*. ACM, 157–168.
- [11] R. Ehlers. 2017. Formal Verification of Piece-Wise Linear Feed-Forward Neural Networks. In *Proceedings of the 15th International Symposium on Automated Technology for Verification and Analysis (ATVA17) (Lecture Notes in Computer Science)*, Vol. 10482. Springer, 269–286.
- [12] E. A. Emerson, A. K. Mok, A. P. Sistla, and J. Srinivasan. 1992. Quantitative temporal reasoning. *Real-Time Systems* 4, 4 (1992), 331–352.
- [13] P. Gammie and R. van der Meyden. 2004. MCK: Model Checking the Logic of Knowledge. In *Proceedings of 16th International Conference on Computer Aided Verification (CAV04) (Lecture Notes in Computer Science)*, Vol. 3114. Springer, 479–483.
- [14] I. Griva, S. Nash, and A. Sofer. 2009. *Linear and nonlinear optimization*. Vol. 108. Siam.
- [15] Z. Gu, E. Rothberg, and R. Bixby. 2016. Gurobi Optimizer Reference Manual. <http://www.gurobi.com>. (2016).
- [16] S. S. Haykin. 1999. *Neural Networks: A Comprehensive Foundation*. Prentice Hall.
- [17] C. Huang, J. Fan, W. Li, X. Chen, and Q. Zhu. 2019. ReachNN: Reachability Analysis of Neural-Network Controlled Systems. *ACM Transactions on Embedded Computing Systems (TECS)* 18, 106 (2019), 1–22.
- [18] K. J. Hunt, D. Sbarbaro, R. Zbikowski, and P. J. Gawthrop. 1992. Neural networks for control systems—A survey. *Automatica* 28, 6 (1992), 1083–1112.
- [19] R. Ivanov, J. Weimer, R. Alur, G. J. Pappas, and I. Lee. 2019. Verisig: verifying safety properties of hybrid systems with neural network controllers. In *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control (HSCC19)*. 169–178.
- [20] K. Julian, J. Lopez, J. Brush, M. Owen, and M. Kochenderfer. 2016. Policy Compression for Aircraft Collision Avoidance Systems. In *Proceedings of the 35th Digital Avionics Systems Conference (DASC16)*. IEEE, 1–10.
- [21] K. D. Julian and M. J. Kochenderfer. 2019. A Reachability Method for Verifying Dynamical Systems with Deep Neural Network Controllers. *CoRR abs/1903.00520* (2019).
- [22] G. Katz, C. W. Barrett, D. L. Dill, K. Julian, and M. J. Kochenderfer. 2017. Reluplex: An Efficient SMT Solver for Verifying Deep Neural Networks. In *Proceedings of the 29th International Conference on Computer Aided Verification (CAV17) (Lecture Notes in Computer Science)*, Vol. 10426. Springer, 97–117.
- [23] P. Kouvaros and A. Lomuscio. 2016. Parameterised Verification for Multi-Agent Systems. *Artificial Intelligence* 234 (2016), 152–189.
- [24] A. Krizhevsky, I. Sutskever, and G. E. Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Proceedings of the 26th Conference on Neural Information Processing Systems (NIPS12)*. Curran Associates, Inc., 1097–1105.
- [25] A. Lomuscio and L. Maganti. 2017. An approach to reachability analysis for feed-forward ReLU neural networks. *CoRR abs/1706.07351* (2017).
- [26] A. Lomuscio, H. Qu, and F. Raimondi. 2017. MCMAS: A Model Checker for the Verification of Multi-Agent Systems. *Software Tools for Technology Transfer* 19, 1 (2017), 9–30.
- [27] P. Maes. 1993. Modeling adaptive autonomous agents. *Artificial life* 1, 1–2 (1993), 135–162.
- [28] V. Nair and G. E. Hinton. 2010. Rectified Linear Units Improve Restricted Boltzmann Machines. In *Proceedings of the 27th International Conference on Machine Learning (ICML10)*. Omnipress, 807–814.
- [29] NANESVerify. 2020. Neural Agent operating on a Non-deterministic Environment System Verify, <https://vas.doc.ic.ac.uk/software/neural/>. (2020).
- [30] N. Narodytska. 2018. Formal Analysis of Deep Binarized Neural Networks. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence, (IJCAI18)*. IJCAI, 5692–5696.
- [31] W. Penczek and A. Lomuscio. 2003. Verifying Epistemic Properties of multi-agent systems via bounded model checking. In *Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multi-agent systems (AAMAS03)*. IFAAMAS, 209–216.
- [32] W. Penczek, B. Woźna, and A. Zbrzezny. 2002. Bounded Model Checking for the Universal Fragment of CTL. *Fundamenta Informaticae* 51, 1-2 (2002), 135–156.
- [33] J. Redmon, S. K. Divvala, R. B. Girshick, and Ali. Farhadi. 2016. You Only Look Once: Unified, Real-Time Object Detection. *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR16)* (2016), 779–788.
- [34] R. S. Sutton and A. G. Barto. 1998. *Reinforcement Learning – An Introduction*. MIT Press.
- [35] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. 2014. Intriguing properties of neural networks. In *Proceedings of the 2nd International Conference on Learning Representations (ICLR14)*.
- [36] V. Tjeng, K. Xiao, and R. Tedrake. 2019. Evaluating Robustness of Neural Networks with Mixed Integer Programming. In *Proceedings of the 7th International Conference on Learning Representations (ICLR19)*.
- [37] S. Wang, K. Pei, J. Whitehouse, J. Yang, and S. Jana. 2018. Efficient Formal Safety Analysis of Neural Networks. In *Proceedings of the 32nd Conference on Neural Information Processing Systems (NIPS18)*. Curran Associates, Inc., 6367–6377.
- [38] W. Winston. 1987. *Operations research: applications and algorithms*. Duxbury Press.
- [39] W. Xiang, H. Tran, J. A. Rosenfeld, and T. T. Johnson. 2018. Reachable Set Estimation and Safety Verification for Piecewise Linear Systems with Neural Network Controllers. In *2018 Annual American Control Conference (ACC)*. AACC, 1574–1579.