Learning in automated negotiation

Renting, B.M.

SUMMARY

This dissertation investigates the design and evaluation of automated negotiation agents within the area of multi-agent systems, with an emphasis on the development of learning negotiation agents to improve performance and generalisability across diverse negotiation settings. Traditional approaches to developing negotiation agents rely on manually designed heuristics and strategies. While effective in specific scenarios, such methods suffer from limitations including human development cost, poor generalisation across different negotiation problems and opponent types, and the introduction of human-induced biases in strategy design. The first part of this dissertation addresses these limitations by exploring methods for learning negotiation strategies. We progress through several methods that enable autonomous agents to learn effective negotiation behaviours.

Chapter 3 presents an automated algorithm configuration technique, specifically Sequential Model-based Algorithm Configuration (SMAC) [75], for the optimisation of the parameters of a manually defined negotiation strategy. This approach demonstrates the potential to improve upon manually tuned or literature-based configurations through computational optimisation, while still relying on a predefined, parametrised strategy structure and manually engineered instance features to guide the configuration process.

Recognising that a single optimised strategy may not be universally optimal, Chapter 4 extends this work by exploring portfolio-based approaches. Using methods such as Hydra [162] for portfolio construction and AutoFolio [100] for per-instance algorithm selection, a portfolio of complementary negotiation strategies is automatically generated, and a selector is trained. This allows the agent to dynamically choose the most suitable strategy from its portfolio based on the characteristics of the current negotiation scenario and opponent, leading to improved adaptiveness and overall performance compared to relying on a single best strategy.

To further reduce the reliance on manual design and potentially mitigate human-induced biases more substantially, Chapter 5 investigates an end-to-end reinforcement learning approach. This method employs Proximal Policy Optimisation (PPO) [135] and uses Graph Neural Networks (GNNs) [86] to handle the variable dimensionality in the observation and action spaces of diverse negotiation problems. This allows the negotiation policy to be learned directly from interaction data without explicit feature engineering or manually designed parametrised strategies. This facilitates further mitigation of human-induced biases and generalisation across negotiation problems of varying sizes and complexities, although with remaining challenges in effective adaptation to opponent types.

In the second part of this dissertation, we move beyond the development of learning agents by critically examining the methods used for their evaluation. Chapter 6 presents an extensive empirical analysis, using data from the Automated

Negotiation Agents Competition (ANAC) 2022, which we organised ourselves and specifically challenged participants to develop learning agents. The analysis shows limitations in current evaluation methods. It demonstrates that agent rankings depend on the choice of performance criteria, such as individual utility, social welfare, or game-theoretic equilibria. It also shows that the commonly used average utility metric is sensitive to group composition and cannot handle non-transitive performance relations between agents. The conclusion drawn is that there is no single, universally robust metric for evaluating negotiation agents, particularly for learning agents that exhibit non-stationary behaviour.

Addressing the identified need for clear evaluation criteria and research challenges grounded in negotiation application domains, Chapter 7 proposes multi-agent calendar scheduling as a rich, relevant, and complex real-world application area. This task contains many challenges for the automated negotiation community, potentially guiding future research efforts to push progress beyond the current boundaries in the automated negotiation community.

In summary, this dissertation contributes methods for developing negotiation agents capable of learning and adapting their strategies, pushing towards reduced human bias and increased generalisability. It also provides a critical analysis of evaluation methodologies in automated negotiation, highlighting their shortcomings and advocating for research into evaluation methods, potentially using application areas like calendar scheduling.