

The ability to learn optimal control policies in systems where action space is defined by sentences in natural language would allow many interesting real-world applications such as automatic optimisation of dialogue systems. Text-based games with multiple endings and rewards are a promising platform for this task, since their feedback allows us to employ reinforcement learning techniques to jointly learn text representations and control policies. We present a general text game playing agent, testing its generalisation and transfer learning performance and showing its ability to play multiple games at once. We also present pyfiction, an open-source library for universal access to different text games that could, together with our agent that implements its interface, serve as a baseline for future research.