

Deep reinforcement learning techniques have demonstrated state-of-the-art performance on board games, which can be represented as sequential combinatorial control problems. Many current, long-standing challenges in engineering are approximately governed by partial differential equation models (e.g., diffusion, electromagnetism, elasticity, options pricing) and can be reduced to combinatorial control problems. We present an algorithm framework that combines a generalized k-opt heuristic with recent advances in deep reinforcement learning. Our results demonstrate the efficacy of deep reinforcement learning as a method for partial differential equation-based optimal control problems with combinatorial constraints. They illustrate the potential of deep reinforcement learning to breathe new life into classical heuristic methods.