Automated planning deals with the problem of finding a sequence of actions leading from a given state to a desired state, e.g., solving Rubik's Cube, delivering parcels, etc. The state-of-the-art automated planning techniques exploit informed forward search guided by a heuristic, where the heuristic estimates a distance from a state to a goal state.

In this thesis, we present a technique to automatically construct an efficient heuristic for a given planning domain.

The proposed approach is based on training a deep neural network using a set of previously solved planning problems from the same domain. We use a novel way of extracting features for states which doesn't depend on usage of existing heuristics. The trained network can be used as a heuristic on any problem from the domain of interest without any limitation on the problem size. Our experiments show that the technique is competitive with popular domain-independent heuristic.

We also introduce a theoretical framework to formally analyze behavior of learned heuristics. We state and prove several theorems that establish bounds on the worst-case performance of learned heuristics.