

This work considers a fundamental open problem in informatics - distinguishing two words by a deterministic finite automaton with the smallest possible number of states. We review the existing research, where proven lower and upper bounds in terms of words' lengths differ exponentially. Next, we empirically try two approaches not studied previously: analysis of discerning sets and application of randomly generated automata. We show that the first approach does not help improve the bounds for the main problem, while the random automata could be successful for randomly taken word pairs but not for all of them. A combination of a randomly generated automaton with the best performing already known, though, may decrease an average number of states by several orders of magnitude. We propose several topics for further investigation based on the obtained experimental results.