

Abstract: BanditFuzz: A Reinforcement-Learning based Performance Fuzzer for SMT Solvers

Joseph Scott, Federico Mora, and Vijay Ganesh

Abstract

In this paper, we present a reinforcement-learning based fuzzing system, BanditFuzz, that zeroes in on the grammatical constructs of well-formed inputs that are the root cause of performance slowdown in SMT solvers. BanditFuzz takes the following as input: a grammar G describing well-formed inputs to a set of distinct solvers (say, a target solver T and a reference solver R) that implement the same specification, and a fuzzing objective (e.g., maximize the relative performance difference between T and R). BanditFuzz outputs a list of grammatical constructs that are ranked in descending order by how likely they are to maximize performance differences between solvers T and R . Using BanditFuzz, we constructed two benchmark suites (with 400 floating-point and 300 string instances) that expose performance issues in all considered solvers, namely, Z3, CVC4, Colibri, MathSAT, Z3seq, and Z3str3. We also performed a comparison of BanditFuzz against random, mutation, and evolutionary fuzzing methods and observed up to a 81% improvement based on PAR-2 scores used in SAT competitions.