Automating SAT Solver Research

markus.iser@kit.edu
jakob.bach@kit.edu

An infinite amount of monkeys produces infinite amounts of code. Eventually, some of this code compiles into programs, and some of these programs are SAT solvers, which means that they correctly solve the NP-complete SAT problem.

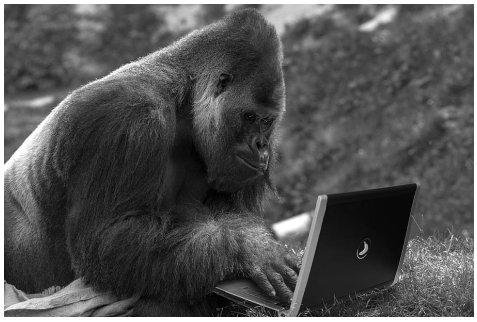

Fig. 1: Monkey writing SAT solver

A research robot incrementally evaluates the efficiency of the SAT solvers which are produced by the monkeys, and deletes those solvers which are not efficient enough. However, it combines the most efficient solvers to a portfolio of solvers [2]. This portfolio can then be used to solve the SAT problem. As new solvers emerge, the robot repeats this process and thus continuously improves its solver portfolio.

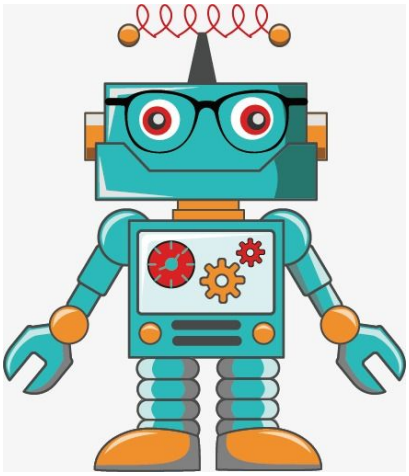

Fig. 2: Research Robot

Whenever a monkey creates a new SAT solver, the research robot determines the efficiency of the solver. For this purpose, the robot measures solver runtimes on SAT instances from a large set I of benchmark instances [1]. The runtime per instance is limited to 2h, such that the robot records either a runtime or a timeout. Based on a scoring function, the robot decides if the solver goes into the portfolio or goes to trash. The scoring function considers the sum of the runtimes and the number of timeouts.

As you can imagine, the monkeys mostly produce SAT solvers which do not improve the overall performance of the solver portfolio. SAT solvers often expose good performance on a subset of instances $S \subseteq I$ and bad performance on the other instances $I \setminus S$. Based on a vector of instance features $v : I \to D^n$, the research robot tries to discriminate the instances of such a subset S from the rest of the instances. If this is possible, the solver is integrated into the portfolio.

However, the research robot has only a limited number of computers in order to determine all the runtimes. Moreover, resource usage comes with a cost. If the research robot gets into debt, the climate changes and all the monkeys die.

You can help us as a human researcher to improve the research robot. We want the robot to predict if a given solver improves the portfolio or goes to trash. The robot should make predictions with high confidence, while conducting a minimal amount of solver runtime experiments.

You will work in an interdisciplinary environment between research on deductive AI (SAT Algorithms) and inductive AI (Big Data Analysis). You get the chance to use state-of-the-art prediction models, clustering techniques and feature selection methods. Help us to build a cost-efficient research robot and save the monkeys!

Fig. 3: Wanted: Human Researcher

References

[1] Iser, M., Springer, L., Sinz, C.: Collaborative management of benchmark instances and their attributes. CoRR **abs/2009.02995** (2020), https://arxiv.org/abs/2009.02995

[2] Xu, L., Hutter, F., Hoos, H.H., Leyton-Brown, K.: Satzilla: Portfolio-based algorithm selection for SAT. CoRR **abs/1111.2249** (2011), http://arxiv.org/abs/1111.2249