

Figure 1: Hardness as Core Refinement Progresses. We run 5 Generations and Core-Refinements for one: Easy problem (left), Hard problem (right).




Figure 2: Left: solver rankings histogram for original SAT Competition Data. Middle: Generated data. We note a striking similarity between the two ranking histograms. Right: per-solver runtime boxplots. We note a similarity between all solvers except 1 and 6, which show discrepancies in the upper quartiles. The medians for these solvers, however, are still similar between original and HardCore-generated problems. We cut off the top of the plot for legibility.

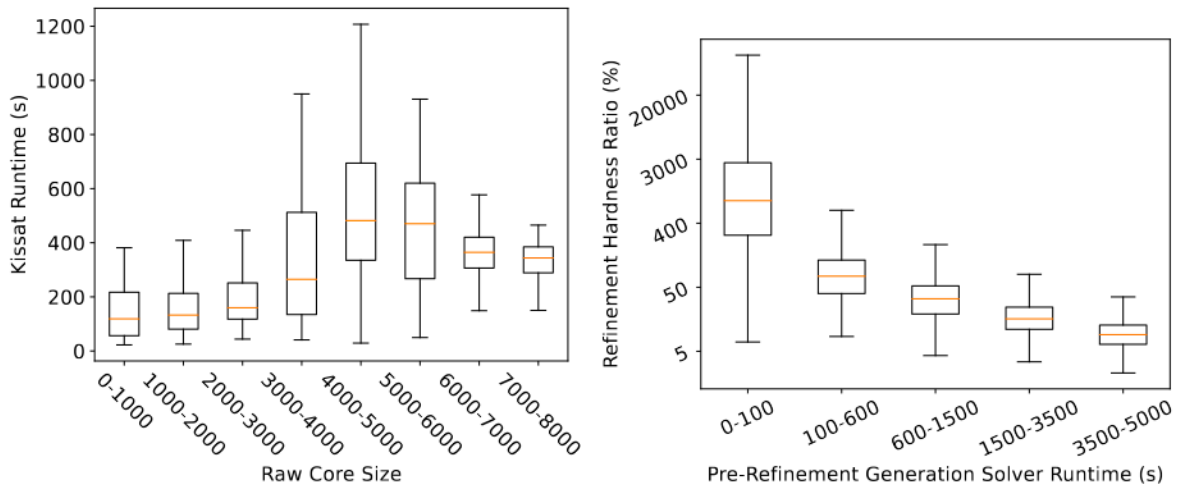


Figure 3: Left: Hardness for problems of varying core sizes. We note a positive correlation between core size and hardness until core size reaches approximately 5000 clauses in size. Right: Post-to-Pre refinement problem hardness ratios, binned by pre-refinement hardness. We see that core refinement recovers the hardness for trivial (hardness-collapsed) problems.