

The Results of SAT Competition 2026

Katalin Fazekas, Marijn Heule, **Ashlin Iser**

SAT 2026 Conference, Lisbon (Portugal)

July 23, 2026

Competition Anniversary

Anniversary of SAT Competitions

- 3 competitions in the 90s (1992,1993, 1996)
- 18 SAT Competitions (2002–)
- 5 SAT Races (2006, 2008, 2010, 2015, 2019)
- 1 SAT Challenge (2012)

Goals

- Promotion of SAT solvers and their development
- Compilation of new challenging benchmarks
- Evaluation of current state-of-the-art solvers

Key rules

- Certified results of unsatisfiability using proof logging
 - ▶ Available Checkers: GRAT, DPR-trim, VeriPB, SR
 - ▶ Instance is “not solved” if proof checker times out
- Disqualification of buggy solvers
 - ▶ Producing an incorrect model or an incorrect proof
- Mandatory solver descriptions + open source
- Ranking scheme: PAR-2
 - ▶ Favors solvers that are faster (not only count solved instances)
- BYOB (Bring Your Own Benchmarks)
 - ▶ At most 20 instances per participant are used

This Year


- Benchmark selection script released before submission deadline
 - ▶ Seed derived from hash of submitted instances
 - ▶ Even organizers cannot predict which instances will be selected
- Special Subtracks for AI Generated / AI Tuned Solvers
 - ▶ Need to outperform the best solver in the category to receive an award
- No-Limits track renamed to Experimental Track
- First year sequential track ran on NHR infrastructure at KIT
 - ▶ absence of memory issues that had plagued the previous competition

Competition Summary

Main Track

- Number of Benchmark Submissions: 1028
 - ▶ 855 by 15 participating teams (> 600 instances by a single team)
 - ▶ 173 by four non-participant benchmark providers
- Number of Sequential Solver Submissions: 33
 - ▶ 2 Sequential Solvers in the Experimental Track
 - ▶ 2 Disqualified (buggy models / wrong results)
 - ▶ 12 AI generated / AI tuned Solvers
 - ▶ 17 Solvers in the Main Track
- Number of Parallel Solver Submissions: 12
 - ▶ 10 Solvers in Parallel Track (1 AI generated / AI tuned Solver)
 - ▶ 2 Solvers in Cloud Track submitted, but one disqualified
- One participant in the Cloud Track

Cloud Track




Parallel Track SAT

Winning Solvers

	Solver	Authors	PAR-2	Solved
AI	hypr-evolve	Jiongzhi Zheng et al.	340.434	145
1	PL-Kissat	Mazigh Saoudi et al.	342.003	148
2	hypr	Jiongzhi Zheng et al.	346.172	146
3	prs-mergesat	Zhenwei Yang et al.	347.677	147

Parallel Track SAT




Parallel Track UNSAT

Winning Solvers

	Solver	Authors	PAR-2	Solved
1	mallob-quick	Dominik Schreiber et al.	182.648	156
2	prs-mergesat	Zhenwei Yang et al.	319.434	135
3	prs-sc26	Zhihan Chen et al.	419.426	119

Parallel Track UNSAT




Parallel Track

Winning Solvers

	Solver	Authors	PAR-2	Solved
1	mallob-quick	Dominik Schreiber et al.	346.504	300
2	prs-mergesat	Zhenwei Yang et al.	401.118	282
3	prs-sc26	Zhihan Chen et al.	453.758	263

Parallel Track




Main Track SAT


Winning Solvers

	Solver	Authors	PAR-2	Solved
AI	Lymphosat	Harrison Green	2813.50	146
1	satsuma-iter+ae-kissat-mab	Markus Anders, Cayden Codel	3430.95	136
2	Kissat-MAB-HyPre	Jiongzhi Zheng et al.	3445.92	134
3	DA-kissat-MAB	Jinghu Liang, Sami Cherif, Chu-Min Li	3545.66	135

Main Track SAT



Main Track SAT




Main Track UNSAT

Winning Solvers

	Solver	Authors	PAR-2	Solved
1	satsuma-iter+kissat	Markus Anders, Cayden Codel	1743.84	149
2	Kissat-MAB-HyPre-V2	Jiongzhi Zheng et al.	3073.61	121
3	EDA-Kissat-v3	Yalun Cai, Qiang Xu	3843.30	111

Main Track UNSAT




Main Track

Winning Solvers

	Solver	Authors	PAR-2	Solved
1	satsuma-iter+kissat	Markus Anders, Cayden Codel	3647.02	276
2	Kissat-MAB-HyPre	Jiongzhi Zheng et al.	3996.29	255
3	EDA-Kissat-v2	Yalun Cai, Qiang Xu	4469.34	240

Main Track



FLoC 2026 Olympic Games

“Open Source Contribution Award”



The award is presented to Armin Biere in recognition of his outstanding contributions to the SAT community through the development of high-performance open-source solvers whose codebases have formed the basis of many SAT Competition winners.

Discussion Points

Lymphosat uses over 100 family-specific LLM-generated solvers using mostly specialized preprocessing (including proofs).

Discussion Point 1: AI Generated / AI Tuned Solvers

- Instance data acts as long-term memory for solvers across runs
 - ▶ Mirrors how humans build on past insight (not inherently a problem)
 - ▶ Natural consequence of generative AI / data-driven heuristics
- Real challenge: measuring generalization beyond the training database, not just memorization
- Possible mitigations
 - ▶ Avoid reusing structurally similar instances in the AI track?
 - ▶ Is it a problem at all if solvers recognize a family?

Discussion Points

This year, we had an unusual number of crashes.

Discussion Point 2: Treatment of Solvers that Crash

Solvers that crash on n (?) instances during execution will be disqualified in the future.

Notably, crashes only occurred in the Experimental and Cloud Tracks.

Recognitions and Thanks

- Katalin Fazekas, Marijn Heule, Ashlin Iser (Organizers)
- Robert Jones, Lee Barnett, Eric Thoman (AWS)
- Tomáš Balyo (Filuta)
- People at NHR@KIT

