A machine learning approach that beats Rubik's cubes

Alexander Chervov[1+], Kirill Khoruzhii[2+], Nikita Bukhal[3], Jalal Naghiyev[2], Vladislav Zamkovoy, Ivan Koltsov[4], Lyudmila Cheldieva[4], Arsenii Sychev[4], Arsenii Lenin[4], Mark Obozov[5], Egor Urvanov[4] and Alexey M. Romanov[4*]

[1] – Institut Curie, Universite PSL;
[2] – Technical University of Munich;
[3] – Novosibirsk State University;
[4] – MIREA-Russian technological university;
[5] – Innopolis University;
[+] – authors contributed equally to this work;
[*] – Corresponding author: romanov@mirea.ru

NeurIPS 2025 accepted Spotlight paper

Moscow, 2025

Rubik’s cube as a benchmarks for AI puzzle sovling

DeepCubeA (2019) and EfficientCube (2023) demonstrated ≈ 70% optimality solving 3×3×3 Rubik’s cube with deep learning approaches.
In 2023, over 1000 teams of ML researchers competed in a Kaggle challenge solving 3×3×3, 4×4×4 and larger cubes.

Can a single unified ML solution beat them all?

Interesting fact: the number of 5×5×5 Rubik’s cube’s states is comparable to the number of atoms in the universe.

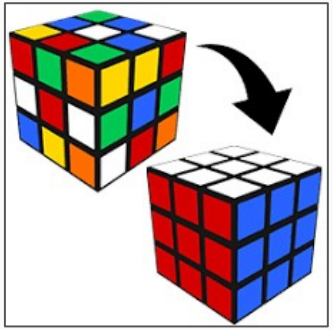

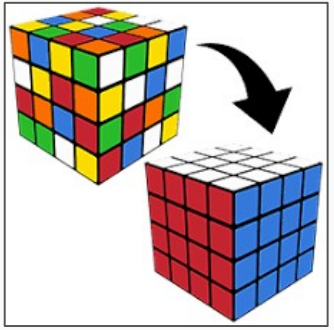

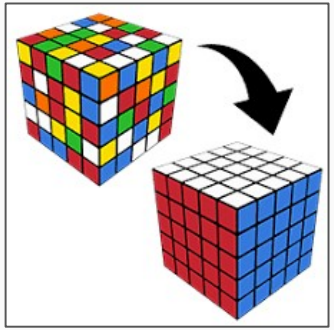

Solution:
(Random Walks + ResMLP + Beam Search) × Agents

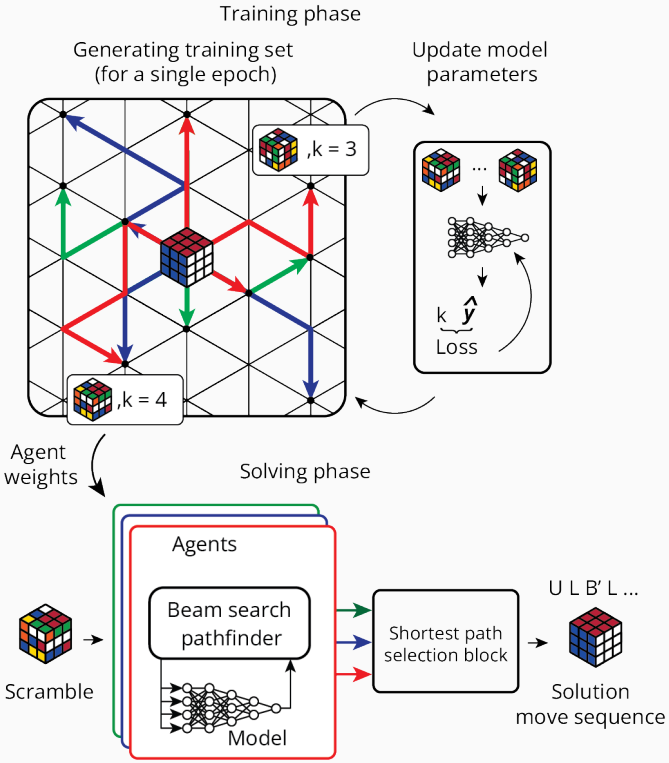

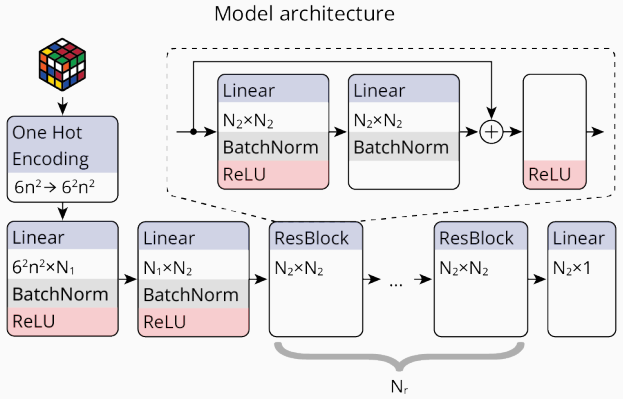

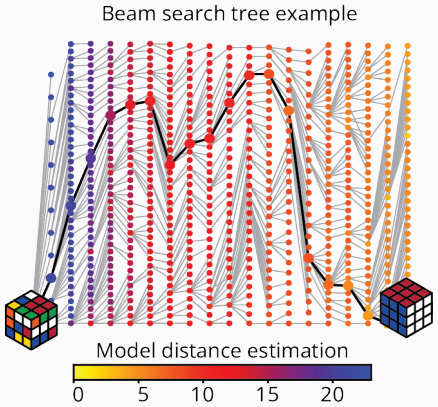

Do not train too much:
we revealed stagnation of solution length

Rubik's Cube 3×3×3

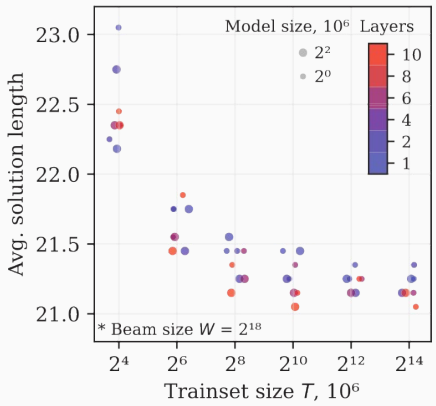

Rubik's Cube 4×4×4

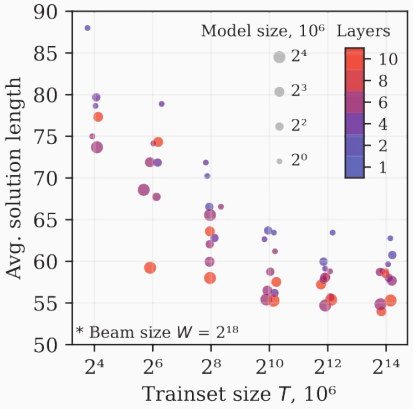

Rubik's Cube 3×3×3

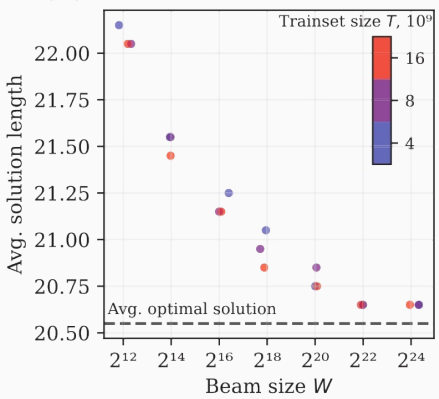

Rubik's Cube 4×4×4

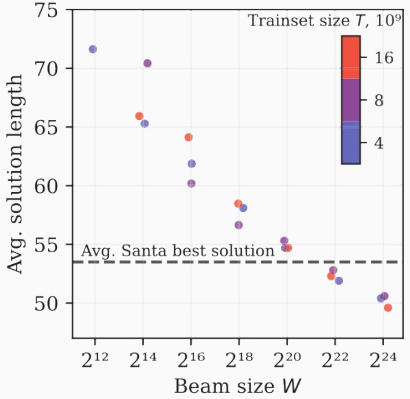

There is no practical reason to use trainsets larger than 8B.

What about really large trainsets? Are you ready for 524B?

Scramble ID

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42

Neural network 1 (upper line 8B/lower line 524B)

49	49	47	45	46	45	51	48	49	47	48	49	50	48	50	47	47	50	48	50	50	50	51	49	48	49	48	48	50	52	49	51	50	50	53	50	50	53	52	51	49	51	53
43	47	45	47	46	49	47	50	47	47	50	45	50	48	40	53	47	52	48	48	50	48	49	47	50	49	48	54	48	48	49	49	50	52	53	44	50	49	50	51	49	51	51

Neural network 2 (upper line 8B/lower line 524B)

51	49	47	45	50	49	51	46	51	49	48	45	50	54	52	47	47	46	48	52	46	50	53	47	48	47	50	50	54	52	49	49	52	52	49	46	58	⨯	52	53	51	51	53
43	47	47	47	48	49	47	52	53	47	⨯	45	48	54	40	51	49	48	50	52	50	48	49	55	50	47	50	48	50	⨯	49	49	50	50	57	46	54	51	52	55	49	51	51

Best results over other exeperiments (neural networks trained with 8B examples)

43	47	43	45	46	45	47	46	47	47	46	45	48	46	40	45	47	44	46	46	44	48	49	47	48	45	46	46	46	48	49	47	48	50	47	42	50	47	48	49	47	49	51

No grokking was observed.

Average solution length slightly improved (49.3→48.56 and 49.74→49.46).

One of networks trained with 524B samples solved less scrambles than its 8B version.

Does this improvement really worth 7 days and 17.5 hours of training?

Multi-agent approach

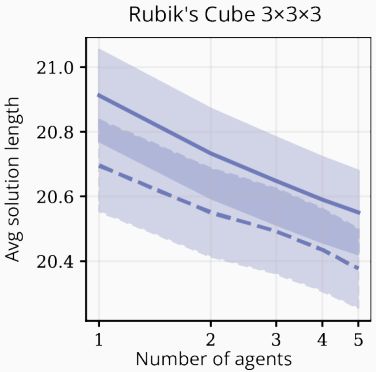

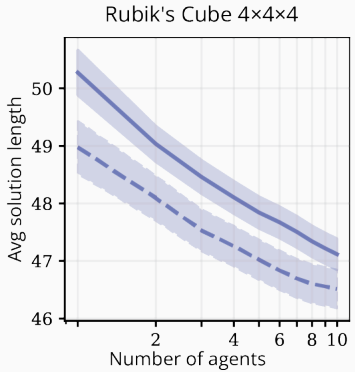

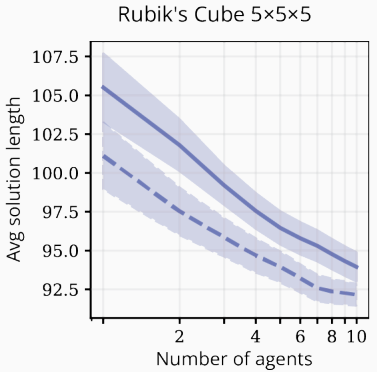

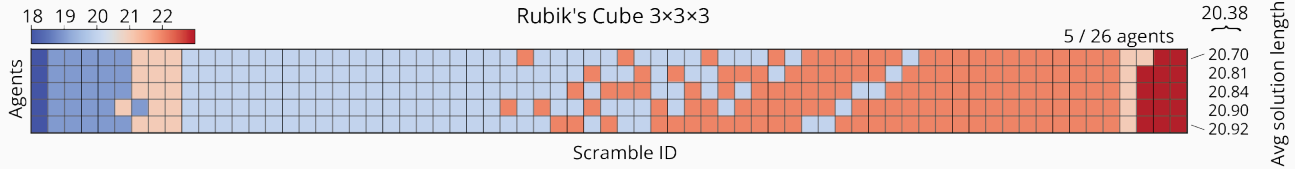

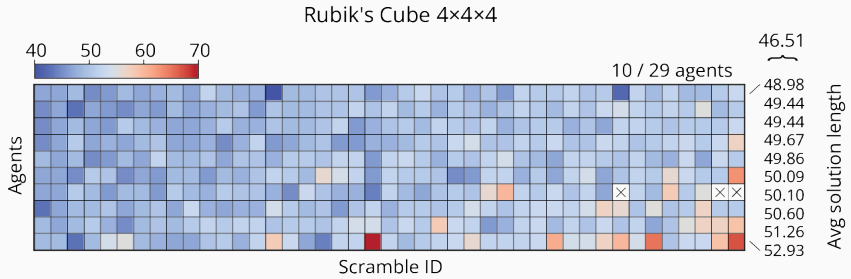

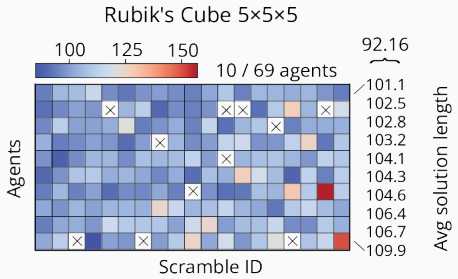

Results

We are significantly faster than competitors!

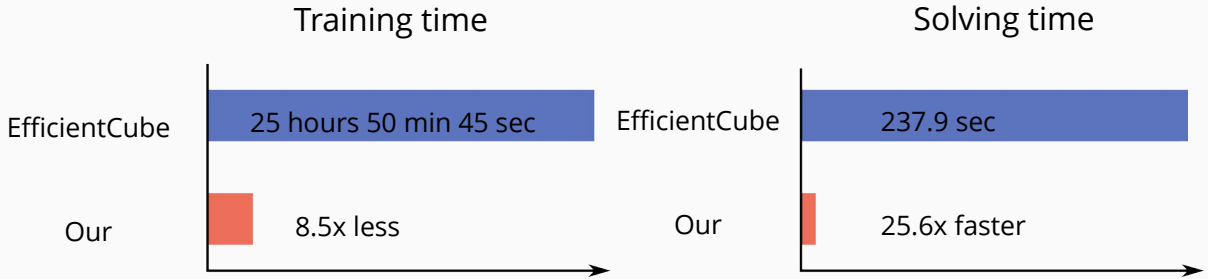

We are more optimal!

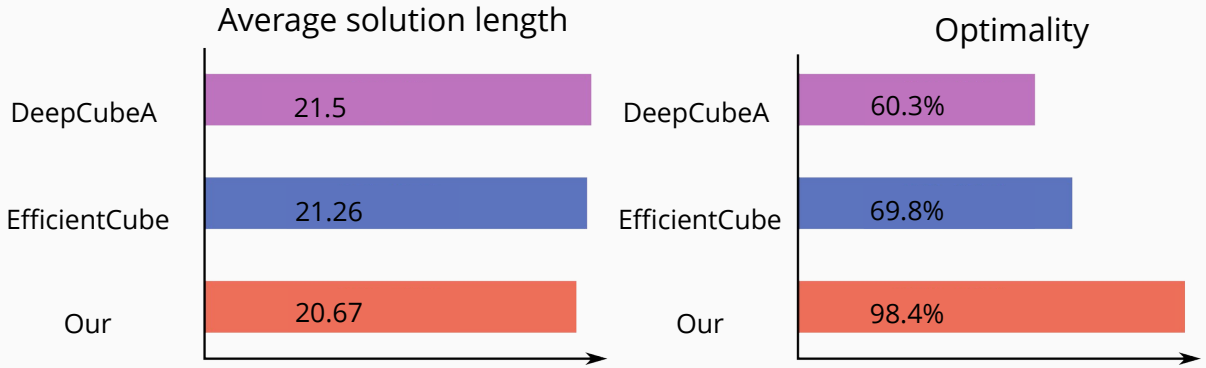

We beat all best 2023 Kaggle Santa Challenge solutions up to 5×5×5 Rubik's cube.

We solved a broad range of puzzles represented by Cayley graphs of size up to 10^{145}.

Thanks!

Contact us:

Check our source

CayleyPy Project

CayleyPy Telegram