

A machine learning approach that beats Rubik's cubes

Alexander Chervov¹⁺, Kirill Khoruzhii²⁺, Nikita Bukhal³, Jalal Naghiyev², Vladislav Zamkovoy,
Ivan Koltsov⁴, Lyudmila Cheldieva⁴, Arsenii Sychev⁴, Arsenii Lenin⁴, Mark Obozov⁵,
Egor Urvanov⁴ and Alexey M. Romanov^{4*}

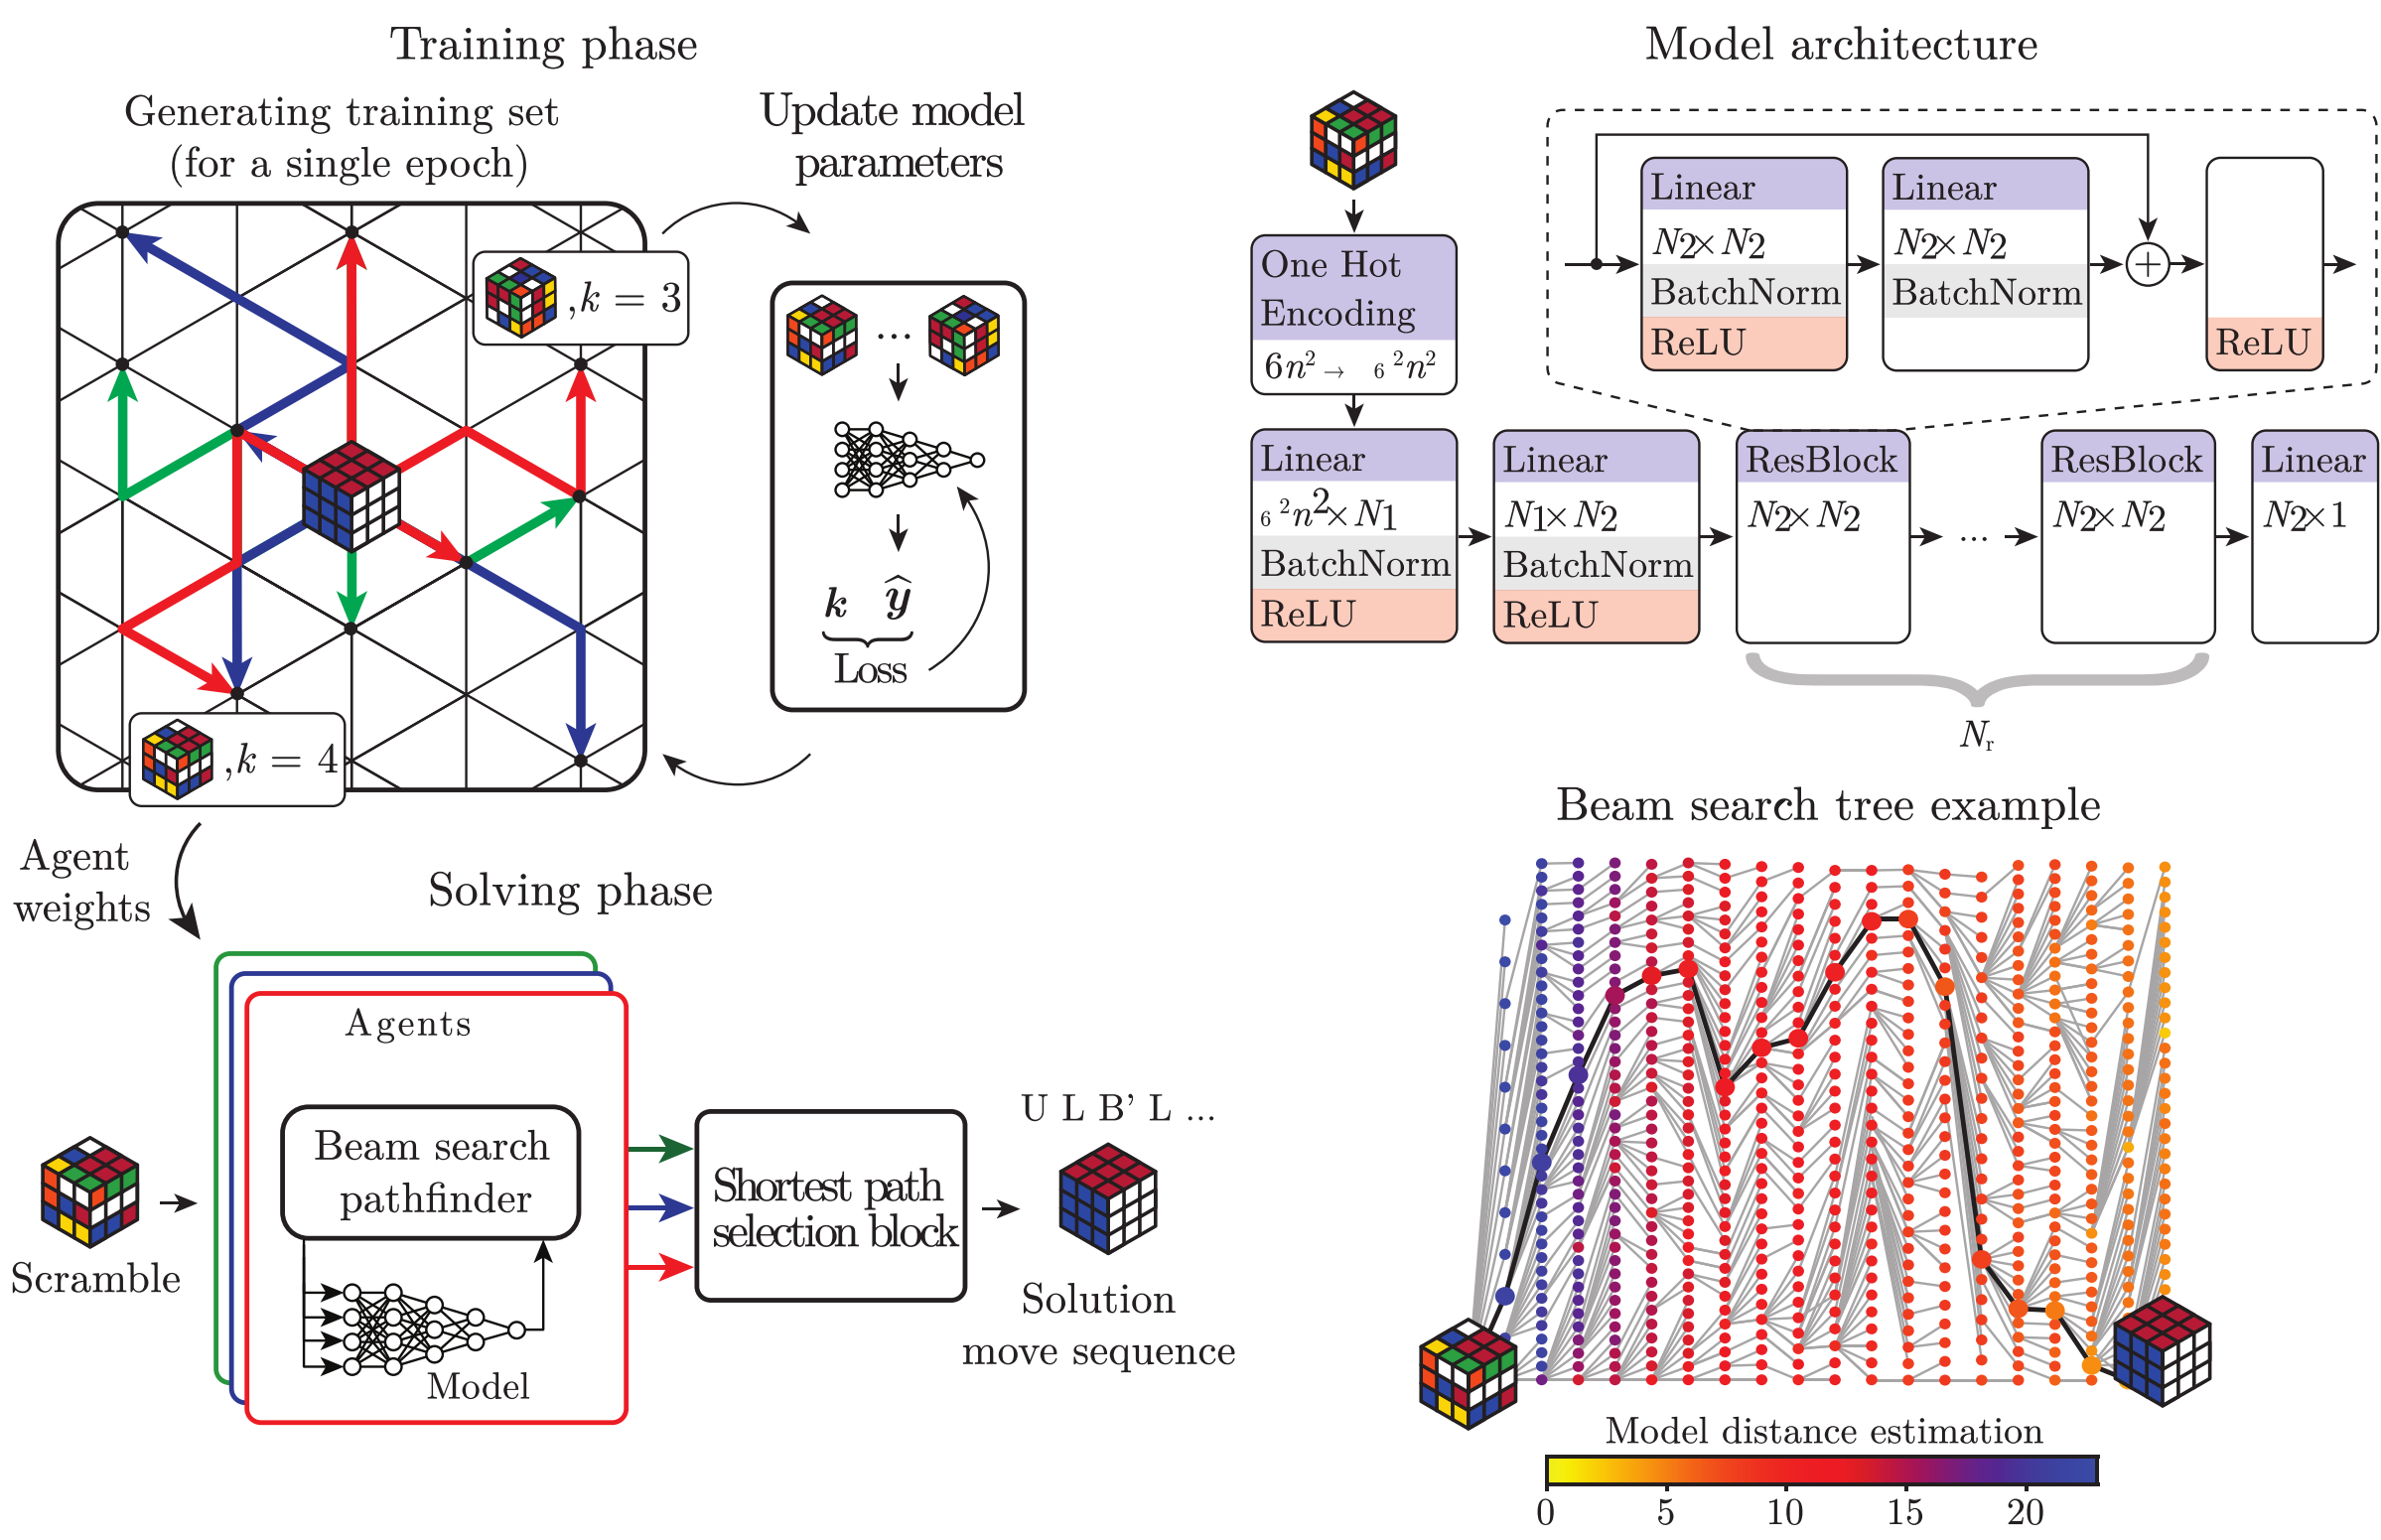
¹ – Institut Curie, Universite PSL, ² – Technical University of Munich, ³ – Novosibirsk State University, ⁴ – MIREA-Russian technological university, ⁵ – Innopolis University
+ – authors contributed equally to this work, * – Corresponding author: romanov@mirea.ru

Can a single unified ML solution beat them all?

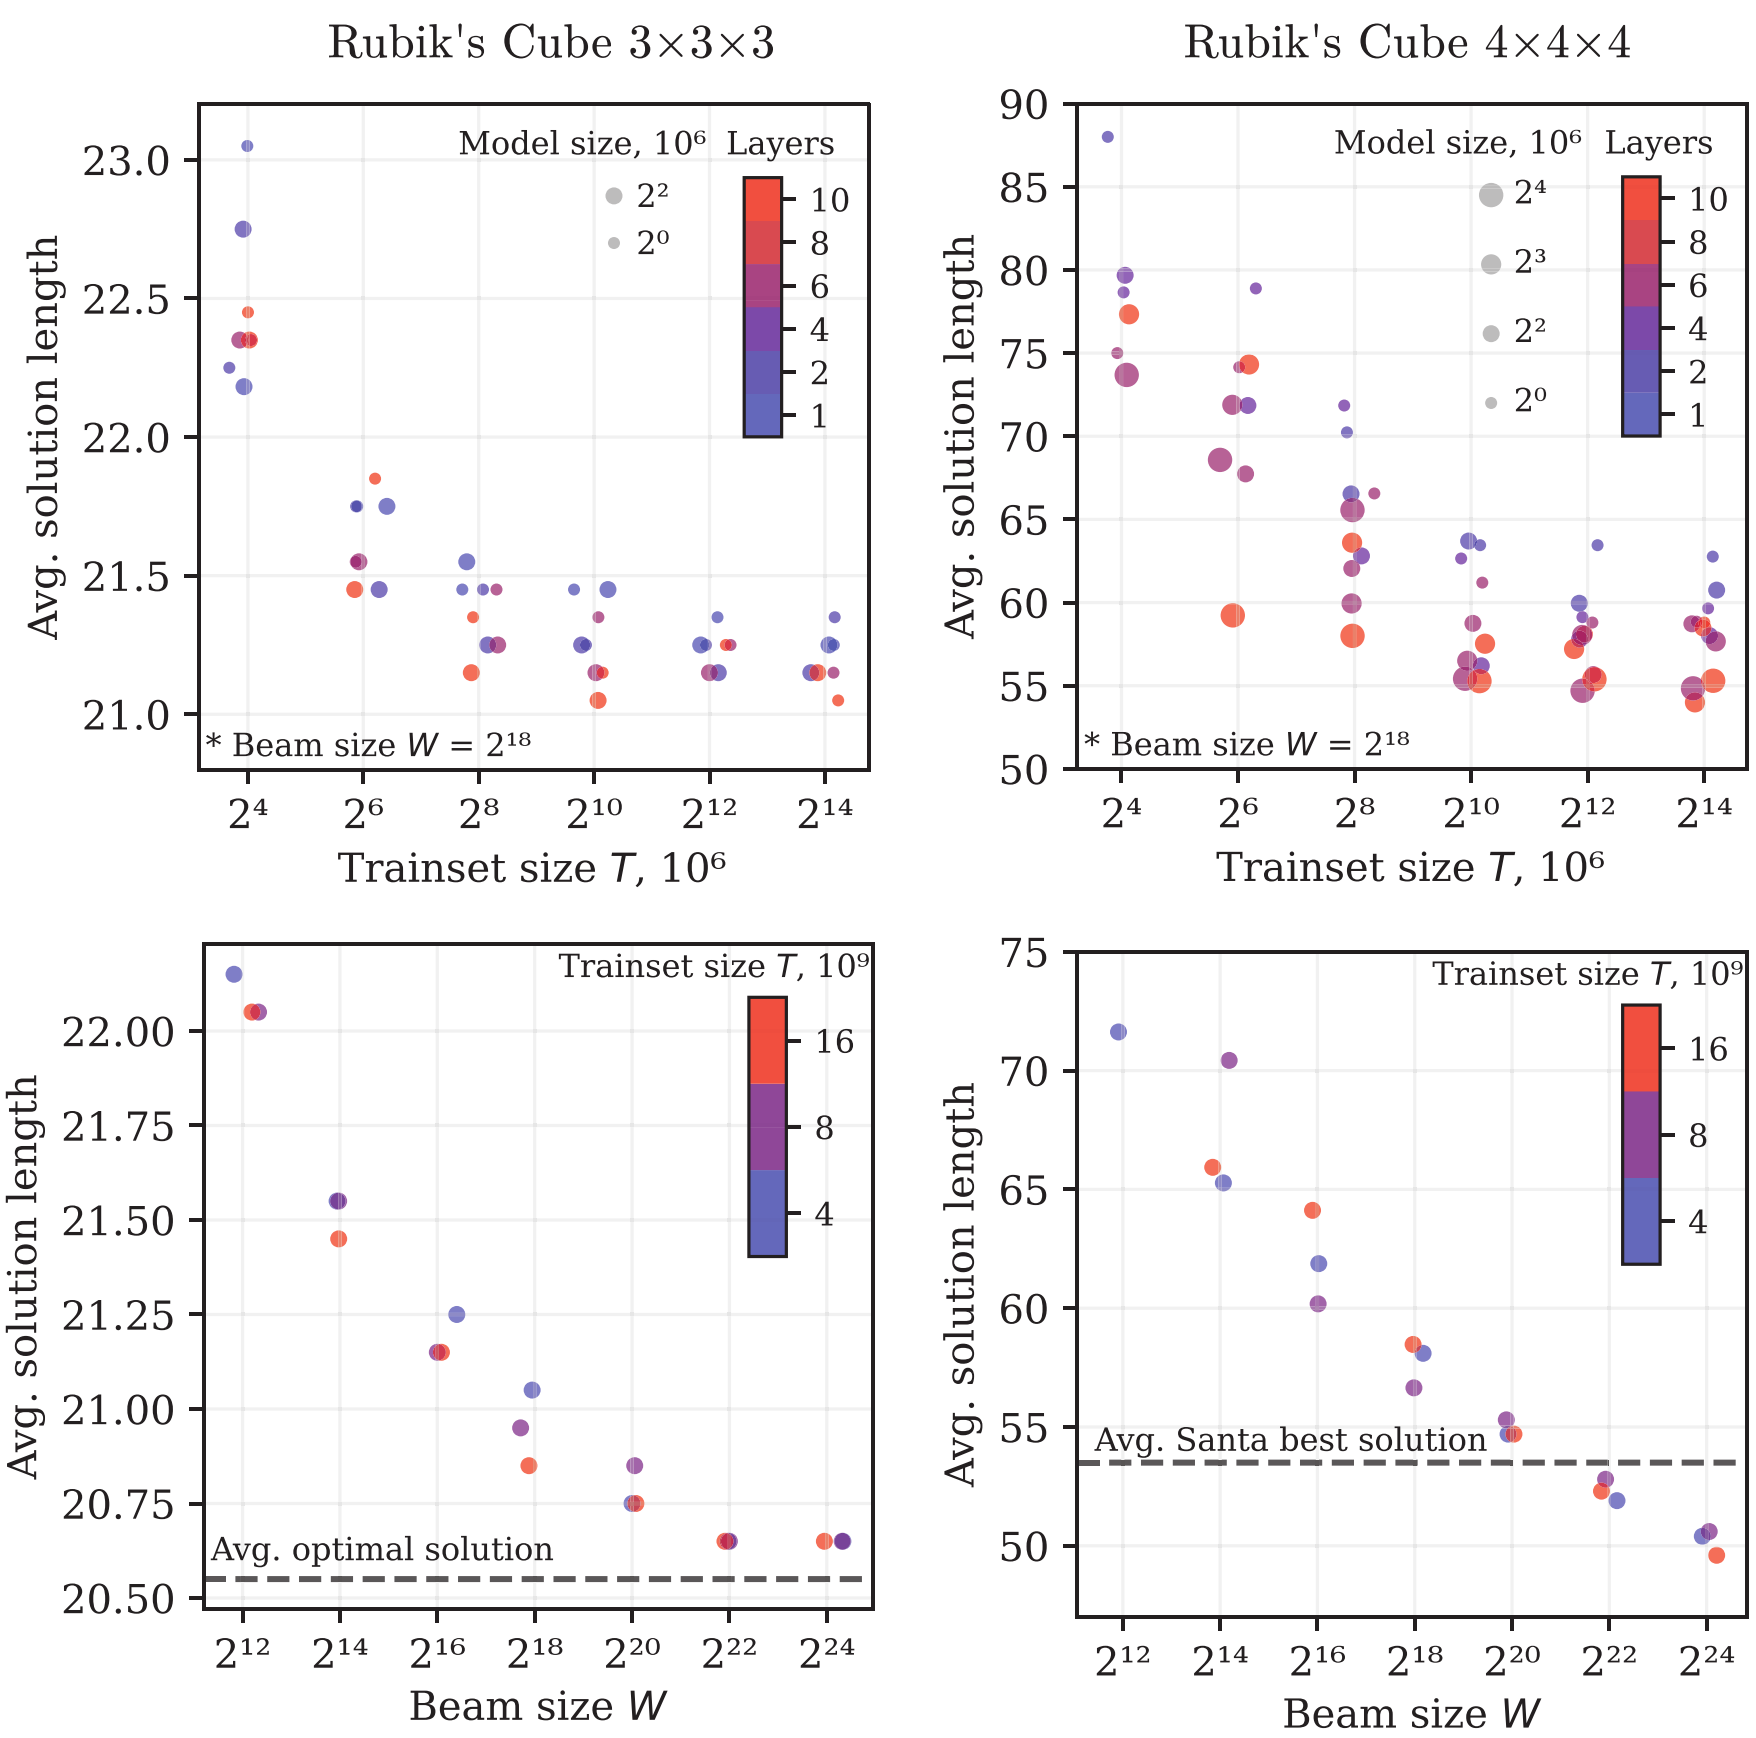
DeepCubeA (2019) and EfficientCube (2023) demonstrated ≈70% optimality solving 3×3×3 Rubik's cube with deep learning approaches.
In 2023, over 1000 teams of ML researchers competed in a Kaggle challenge solving 3×3×3, 4×4×4, and larger cubes.

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

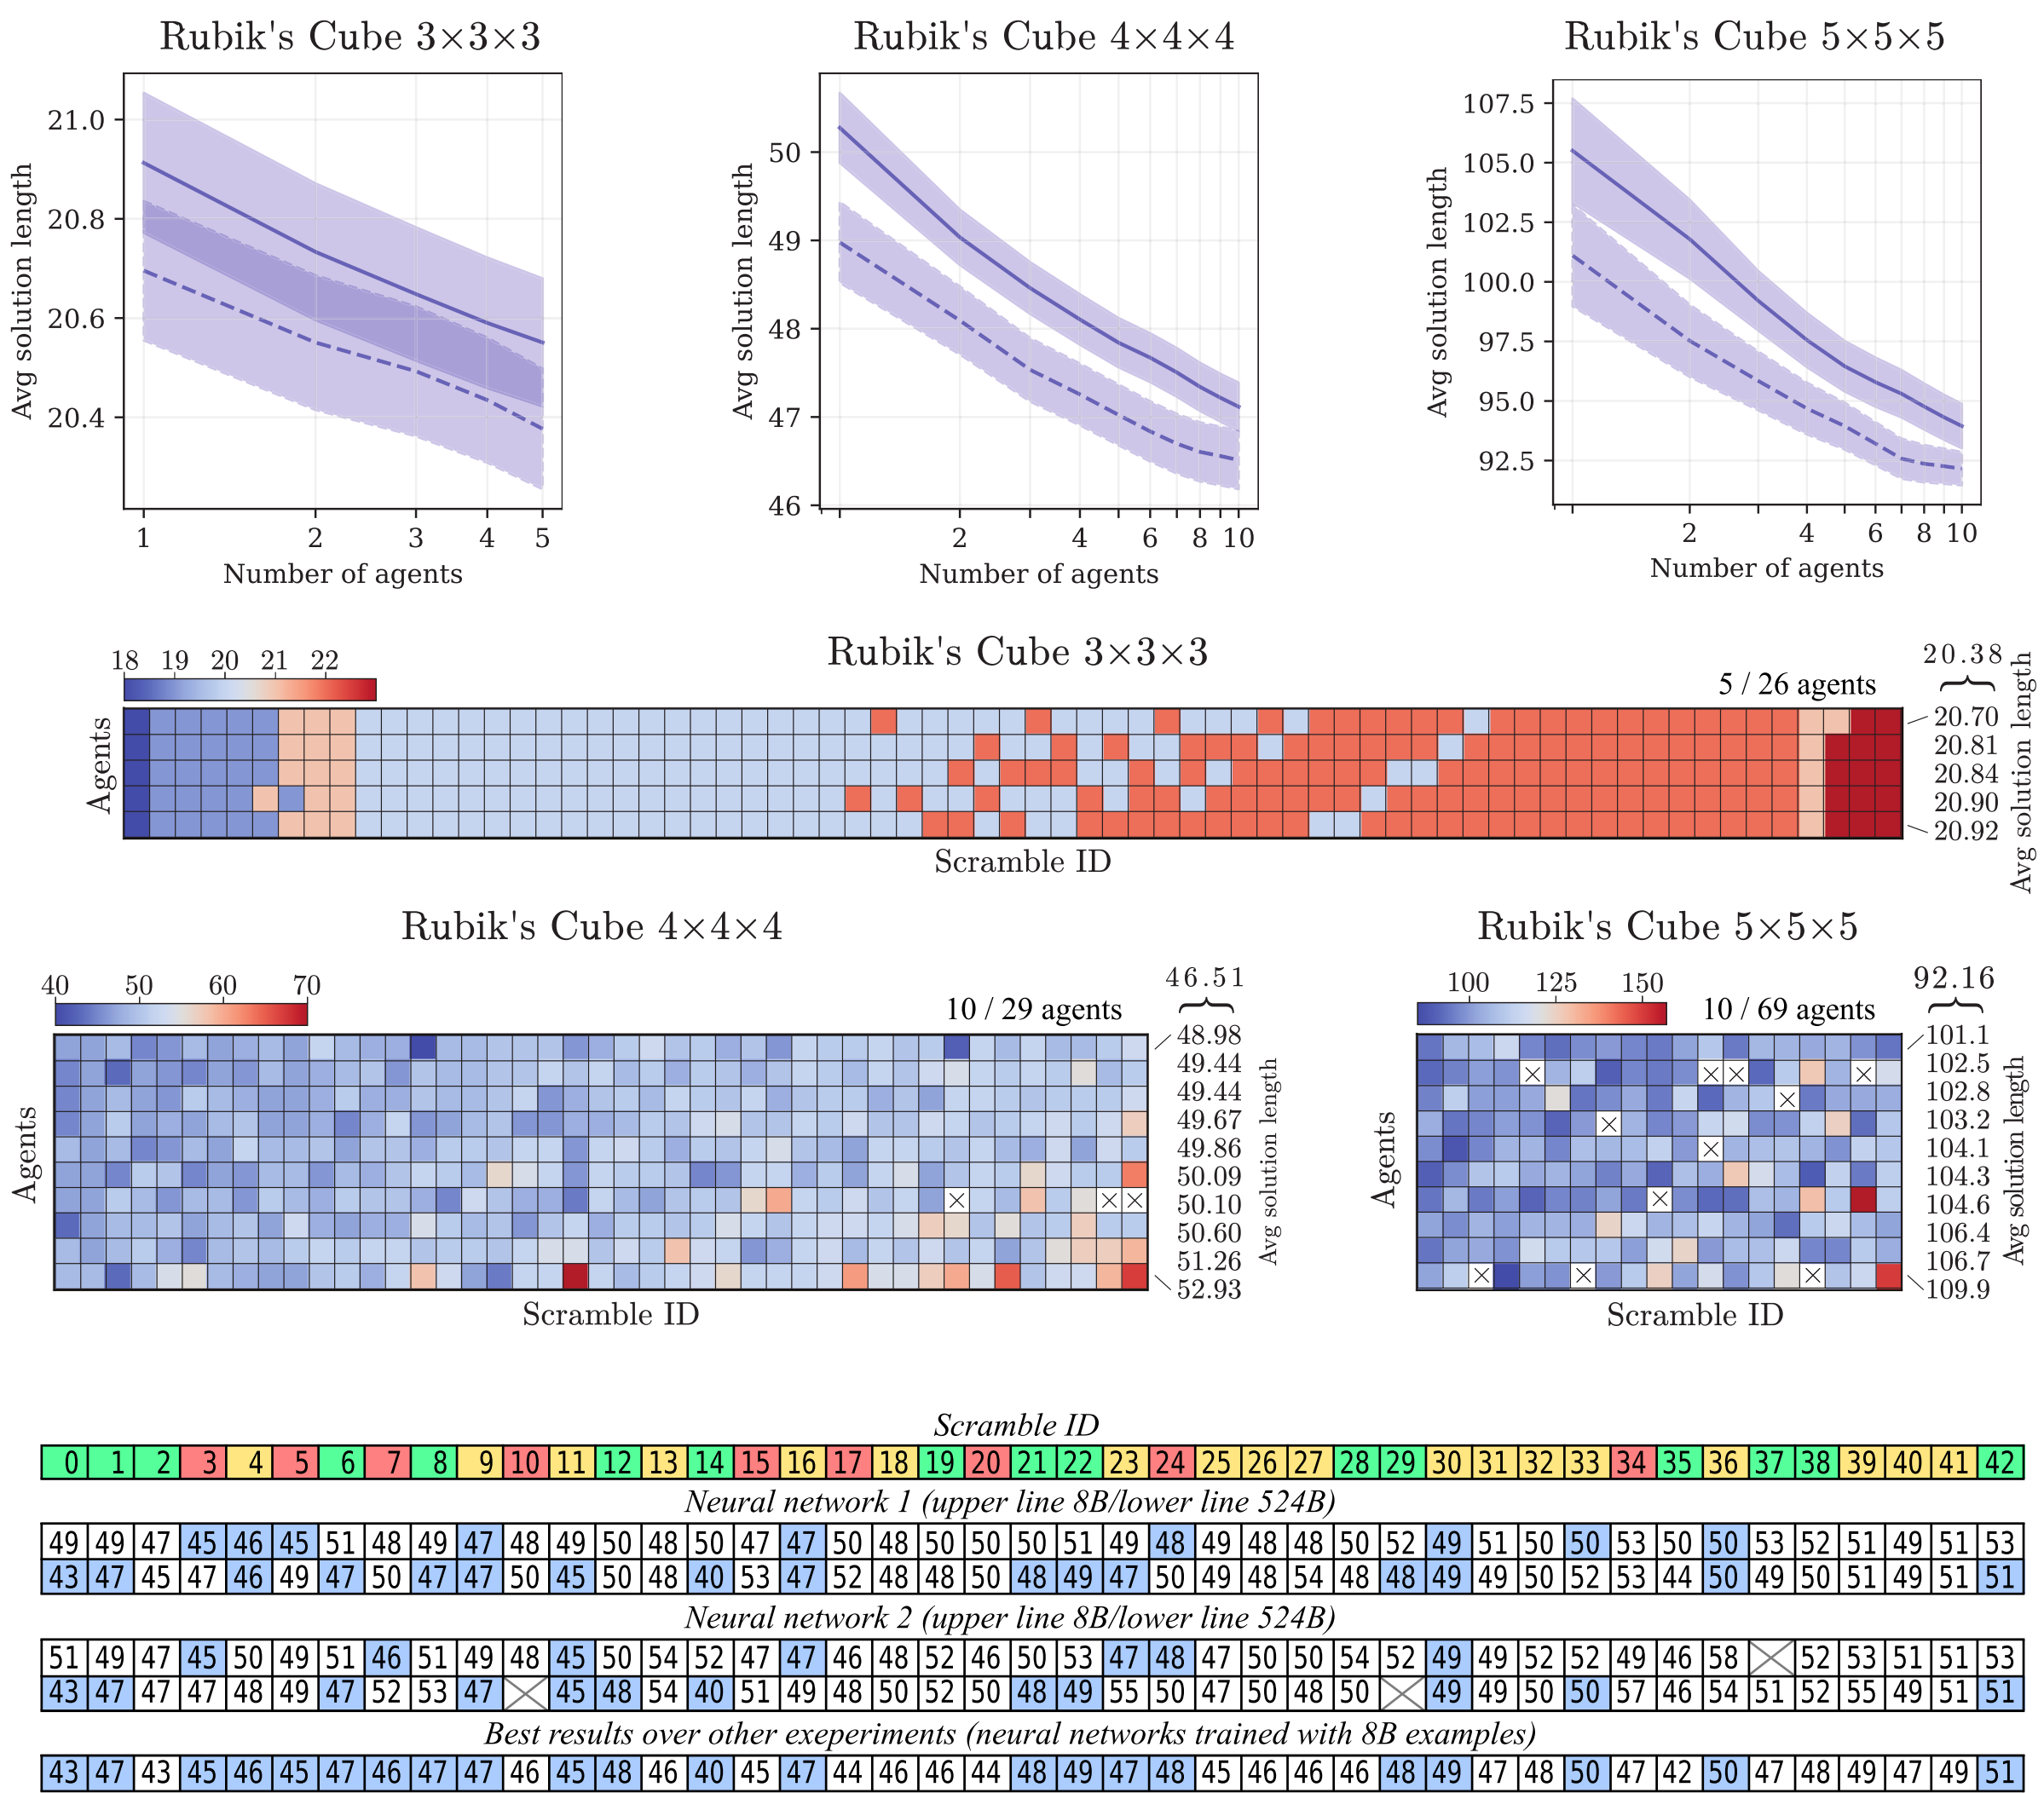
Solution: (Random Walks + ResMLP + Beam Search) × Agents



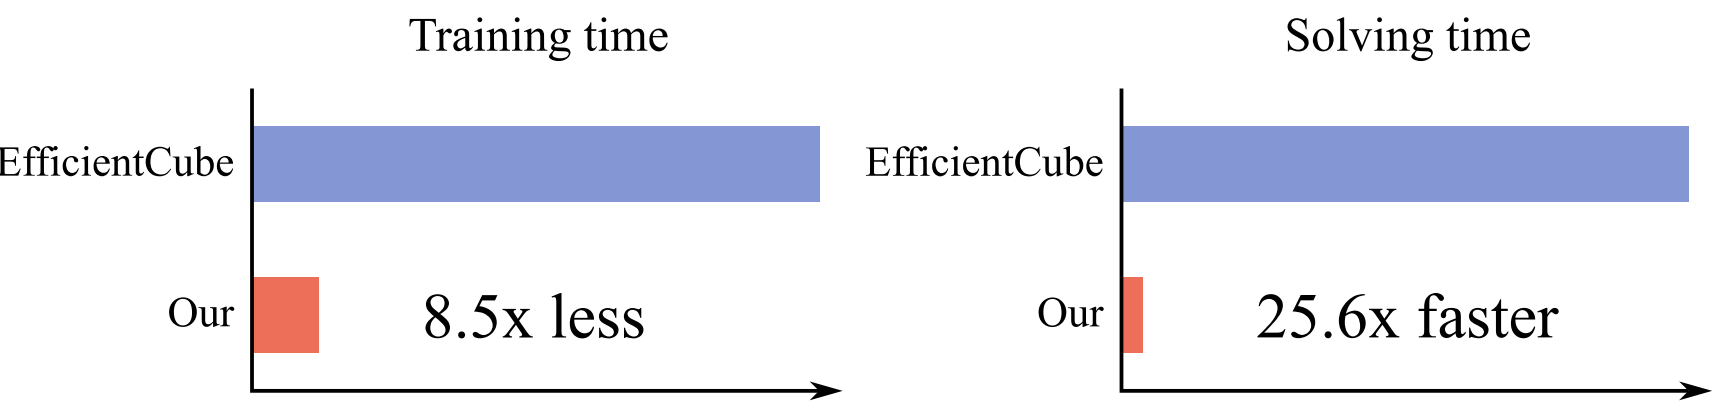
Do not train too much: we revealed stagnation of solution length



Teach on more samples or teach more agents?



Boosting performance:



Results:

- 98.4% solution optimality for 3×3×3 Rubik's cube;
- Beat all best 2023 Kaggle Santa Challenge solutions up to 5×5×5 Rubik's cube;
- Solved a broad range of puzzles represented by Cayley graphs of size up to 10^{145} .

Check our source

