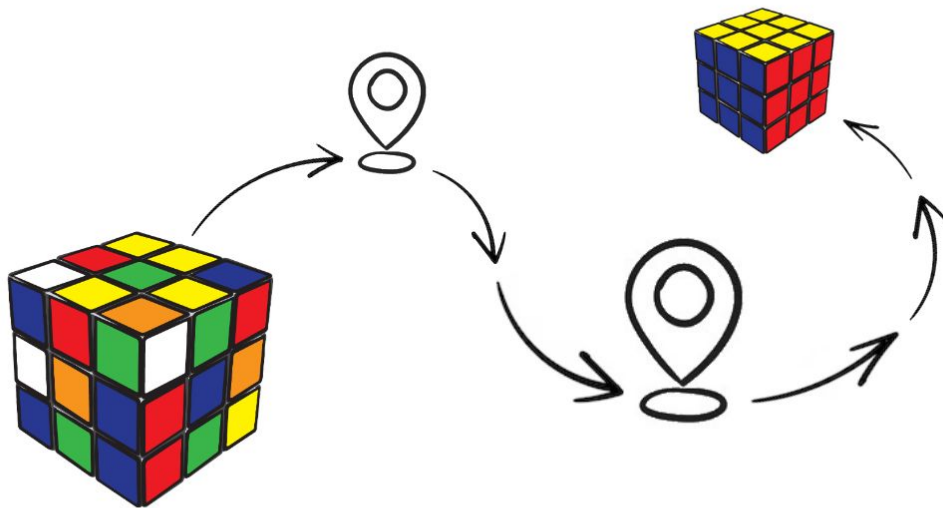




Fast and Precise:
Adjusting Planning Horizon with
Adaptive Subgoal Search



Meet the authors



Michał Zawalski, Michał Tyrolski, Konrad Czechowski, Tomasz Odrzygóźdź,
Damian Stachura, Piotr Piękos, Yuhuai Wu, Łukasz Kuciński, Piotr Miłoś

Environments



INT (proving inequalities)




Sokoban




Rubik's Cube


Graph search



Graph search



Adaptive Subgoal Search



Adaptive Subgoal Search


Adaptive Subgoal Search




Adaptive Subgoal Search




Adaptive Subgoal Search



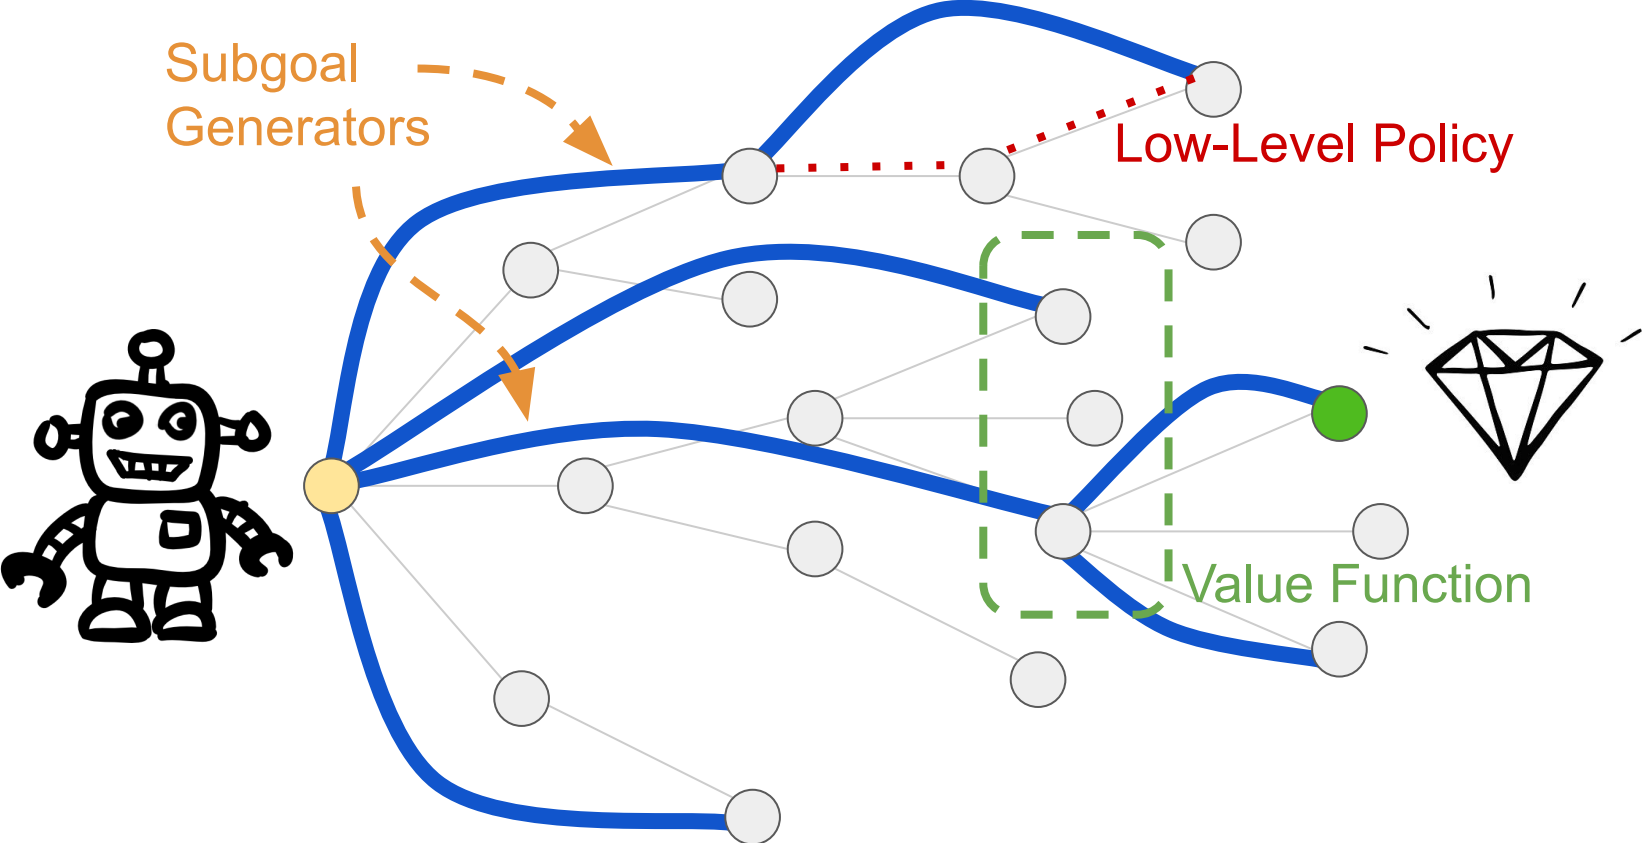
Adaptive Subgoal Search



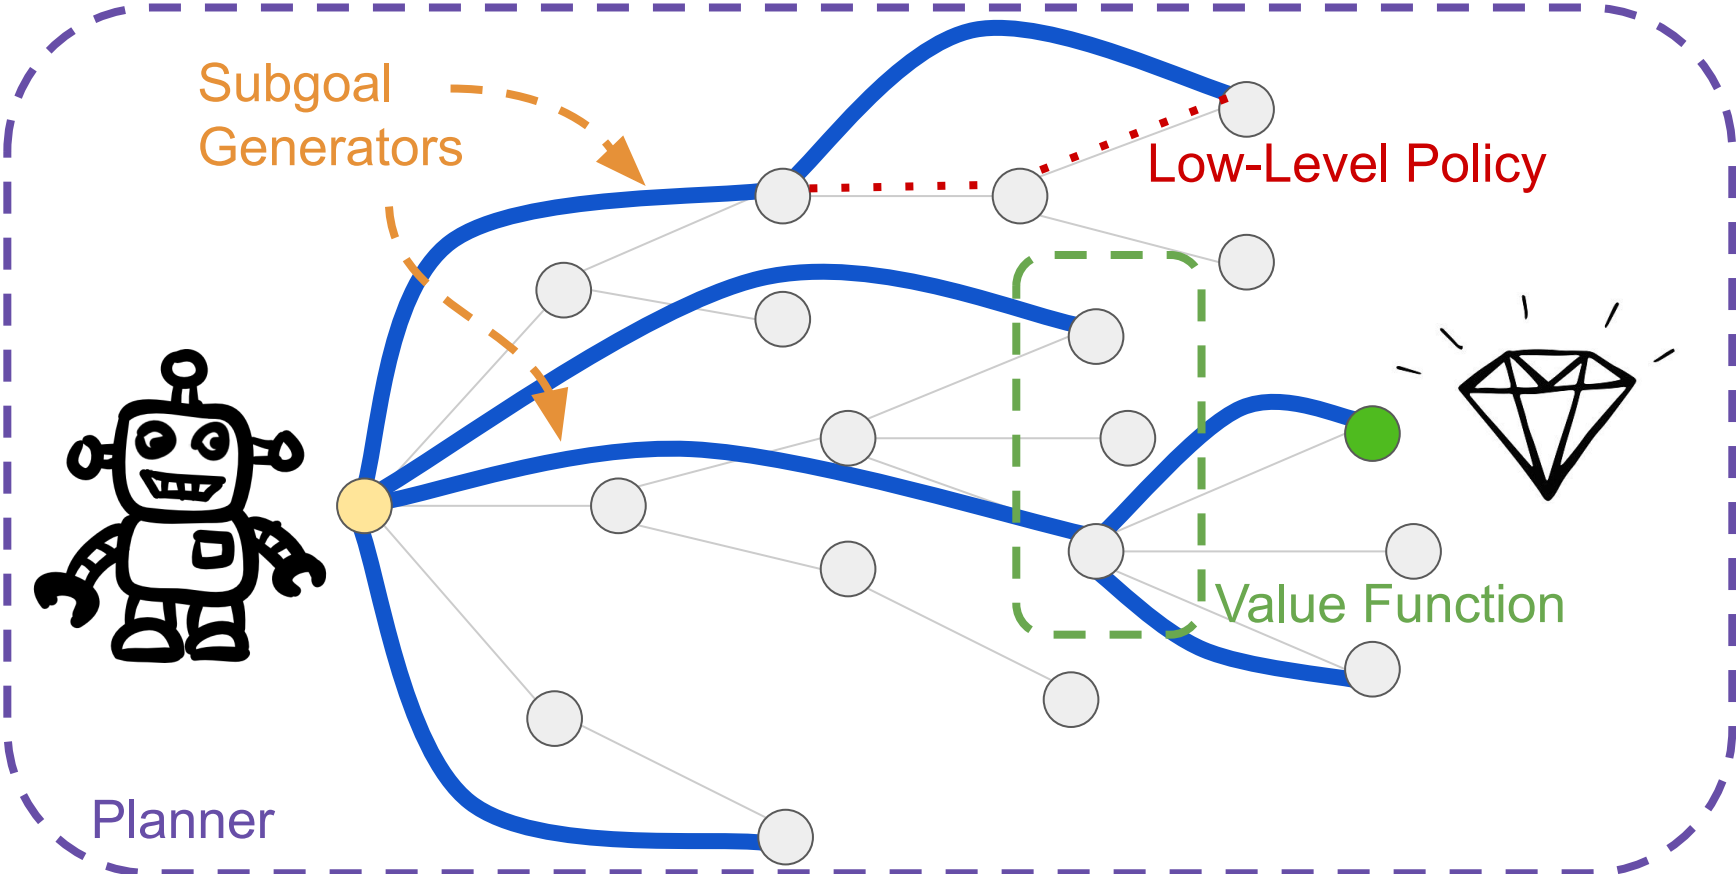
Adaptive Subgoal Search



Adaptive Subgoal Search

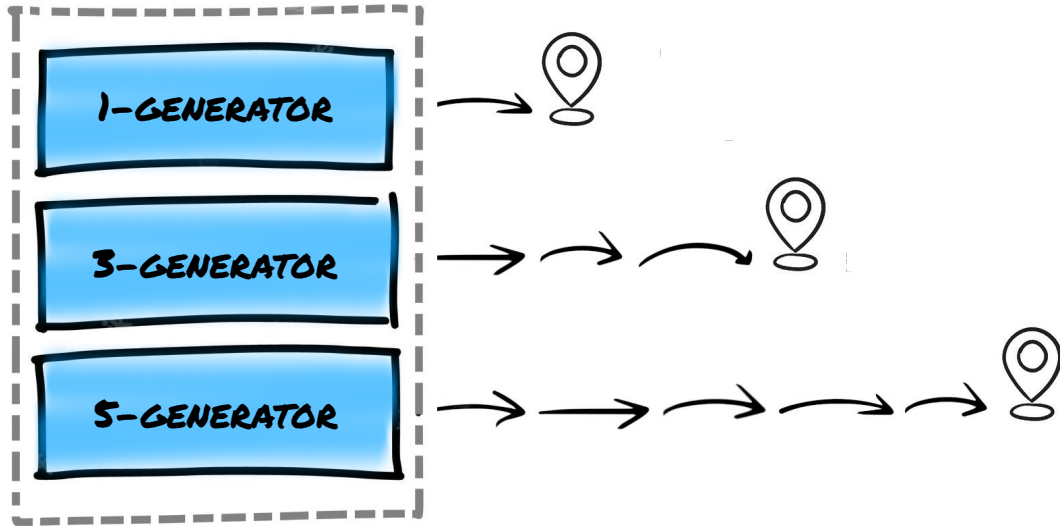


Adaptive Subgoal Search



Developing adaptive search methods

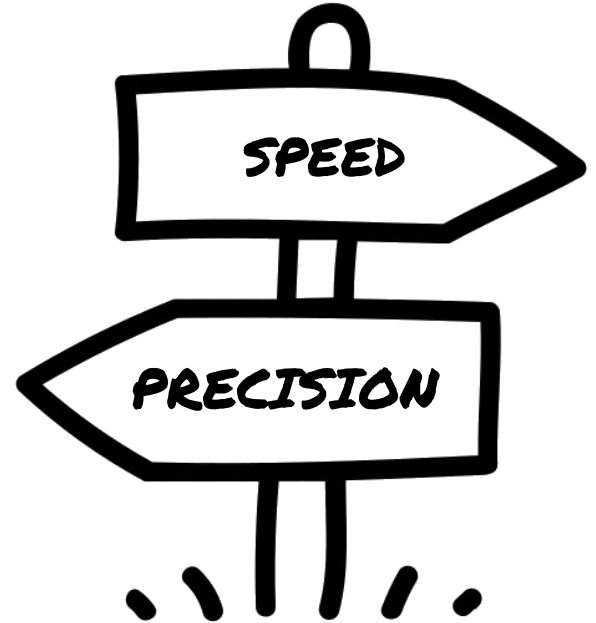
- We achieve adaptivity by using a set of subgoal generators.
- In every step, the generator is chosen by the planner



Developing adaptive search methods

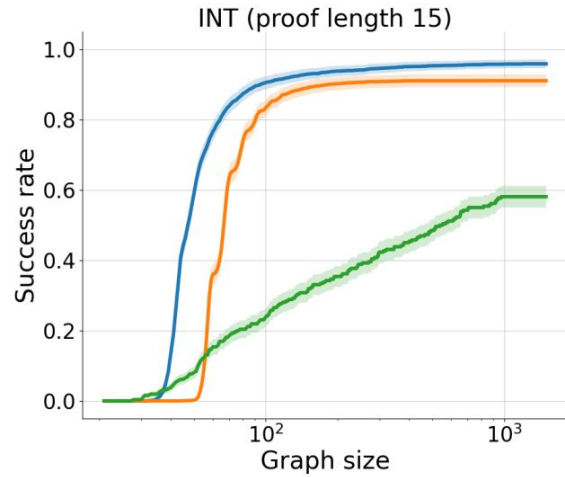
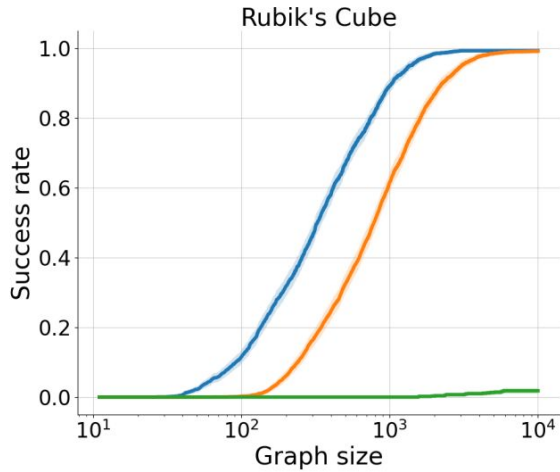
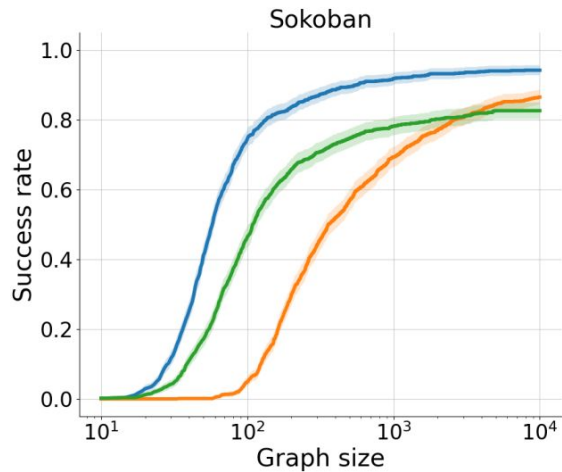
We experimented with several adaptive approaches. Two conclusions follow:

- It is beneficial to strongly prioritize the longest subgoals
- Adaptive methods outperform non-adaptive Subgoal Search

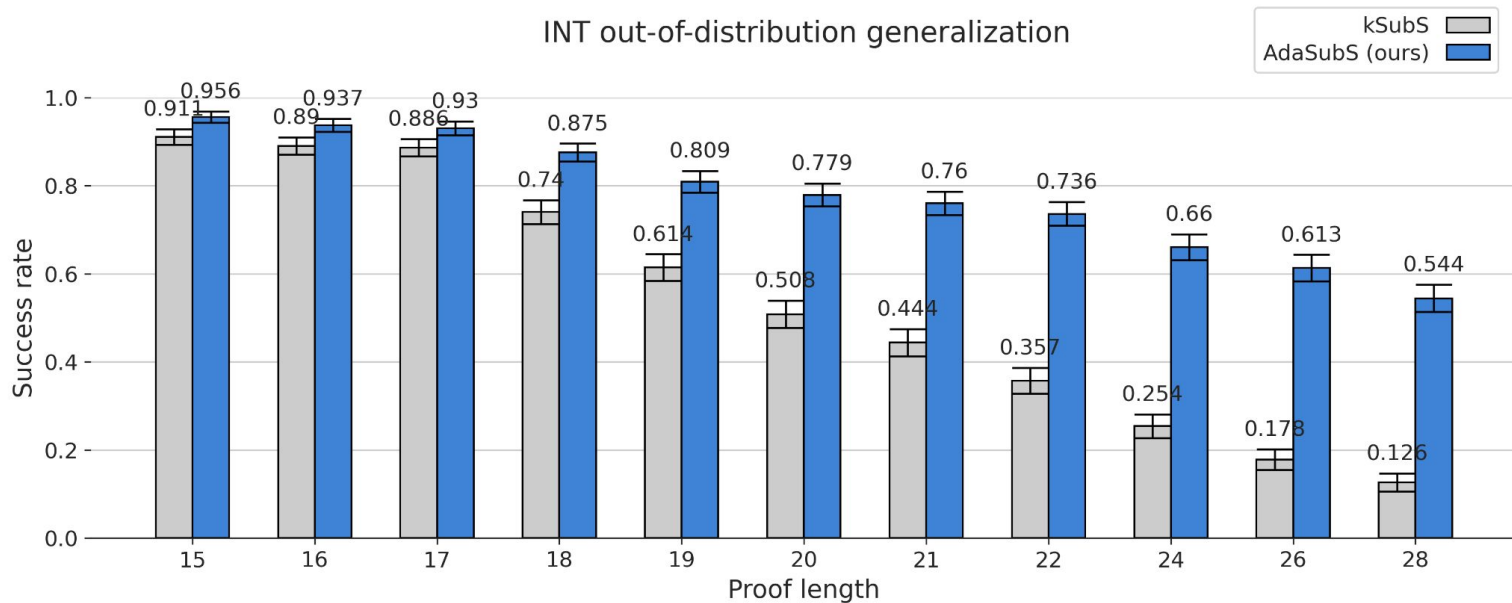


Experimental results


— AdaSubS (ours) — kSubS — BestFS



Out-of-distribution generalization




Why does Adaptive Subgoal Search work?




Why does Adaptive Subgoal Search work?

- decomposes problem into easier parts




Why does Adaptive Subgoal Search work?

- decomposes problem into easier parts
- adapts to the local complexity




Why does Adaptive Subgoal Search work?

- decomposes problem into easier parts
- adapts to the local complexity
- mitigates errors of one component with others



Why does Adaptive Subgoal Search work?

- decomposes problem into easier parts
- adapts to the local complexity
- mitigates errors of one component with others
- uses only supervised training objectives



What remains for research?



What remains for research?



- planning in a latent space

What remains for research?



- planning in a latent space
- model dynamics

What remains for research?



- planning in a latent space
- model dynamics
- stochastic environments

What remains for research?



- planning in a latent space
- model dynamics
- stochastic environments
- solution optimality

What remains for research?




- planning in a latent space
- model dynamics
- stochastic environments
- solution optimality
- real-world applications

The final summary of Adaptive Subgoal Search




Solves complex problems

by decomposing them into subgoals and planning with high-level steps



Adapts to the local complexity

by choosing the appropriate distance for generating subgoals.



Generalizes to out-of-distribution instances

without any additional fine-tuning.

Visit our webpage!

