Finding Increasingly Large Extremal Graphs with AlphaZero and Tabu Search

WPI · Imperial College London · Google DeepMind · Universität Bonn · Yonsei University

Abbas Mehrabian, Ankit Anand, Hyunjik Kim, Nicolas Sonnerat, Matej Balog, Gheorghe Comanici, Tudor Berariu, Andrew Lee, Anian Ruoss, Anna Bulanova, Daniel Toyama, Sam Blackwell, Bernardino Romera Paredes, Petar Veličković, Laurent Orseau, Joonkyung Lee, Anurag Murty Naredla, Doina Precup, Adam Zsolt Wagner

Problem

Given n, find n-node graphs that maximize the number of edges such that the graph has no cycles of size 3 or 4.

f(n) = max number of edges in n-node graph without 3 or 4 cycles

Erdős (1975) conjectured: $\lim_{n\to\infty} \frac{f(n)}{n\sqrt{n}} = \frac{1}{2\sqrt{2}}$

Many lower and upper bounds have been obtained for f(n).

Our goal: improve existing numerical lower bounds for f(n), hoping to find a pattern that could lead to refuting the conjecture.

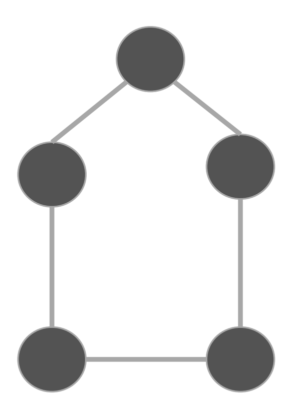

Optimal (max f(n)) graph for n=5

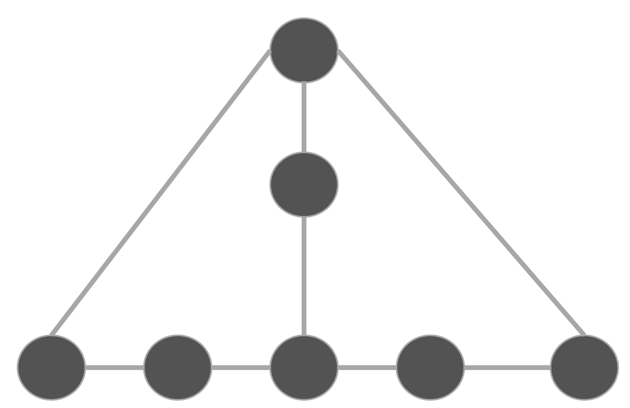

Optimal (max(f(n)) graph for n=7

Graph generation as sequential decision-making

- Transition dynamics and action spaces
  - Start from empty graph and add nodes one by one (short horizon, large action space)
  - Start from empty graph and add edges one by one (large horizon, small action space)
  - Start from a given graph and add/remove edges one by one (moderate horizon ?, moderate action space)
- Rewards
  - Non-telescopic reward
    - 0 rewards elsewhere score(G) at terminal state
  - Telescopic reward

$$r_t = score(G_t) - score(G_{t-1})$$

Methods

AlphaZero

- Widely used method in ML for search problems (AlphaTensor[2022])
- Monte Carlo Tree Search (MCTS) with learned policy and value network
- Carefully balances exploration and exploitation trade-off.
- Distributed implementation with multiple actors and learner
- Joint learning over multiple sizes together in a single run

Tabu Search

- State-of-the-art local search method on many search problems
- Main scheme to improve lower bounds on f(n)
- Bans recently visited states to avoid getting stuck in local minima
- Instead of banning states, we ban recently flipped edges (actions)
- Distributed implementation with many actors

Network Representation

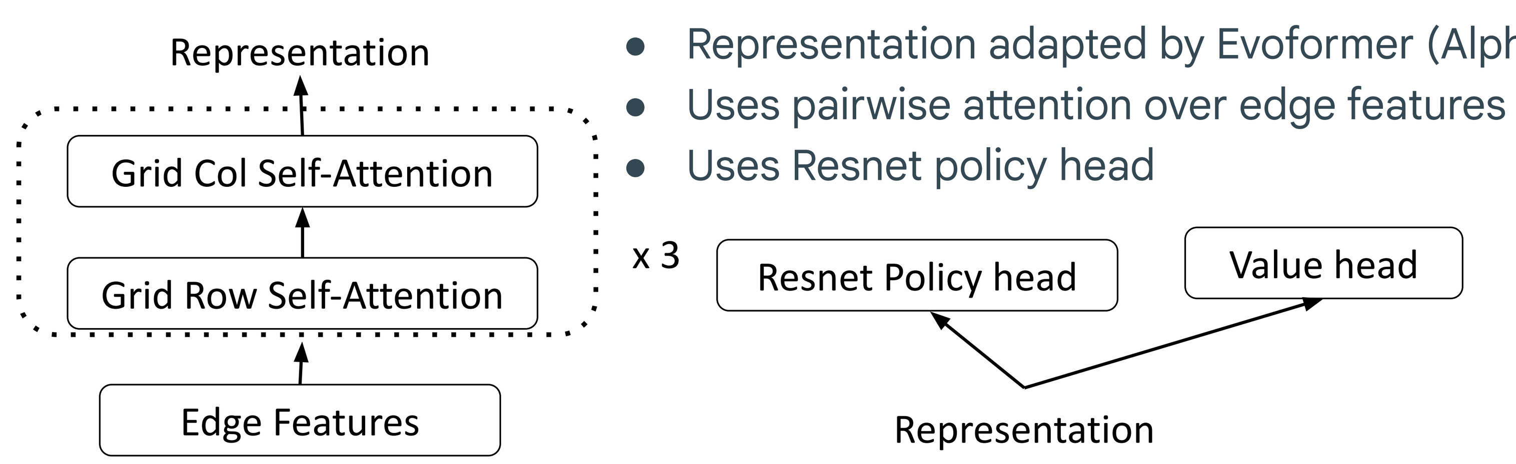

- Representation adapted by Evoformer (AlphaFold)
- Uses pairwise attention over edge features (with layernorm)
- Uses Resnet policy head

Curriculum

Observation:

High-scoring graphs of smaller size are often near-subgraphs of high-scoring graphs of larger sizes.

Leveraging AlphaZero with curriculum

- Use edge-flipping environment
- Generate graphs of smaller size from scratch
- Use graphs generated by smaller size as initial graphs for large size, and iterate

Leveraging incremental tabu search with curriculum

- Change starting graph of tabu search to an already found graph for size (n-k, n-1)

Other implementation details

- For AlphaZero, distributed implementation, run a joint network for multiple sizes where an episode on graph size n samples initial graphs from (n-k, n-1)

Results

Improved lower bounds over the state of the art

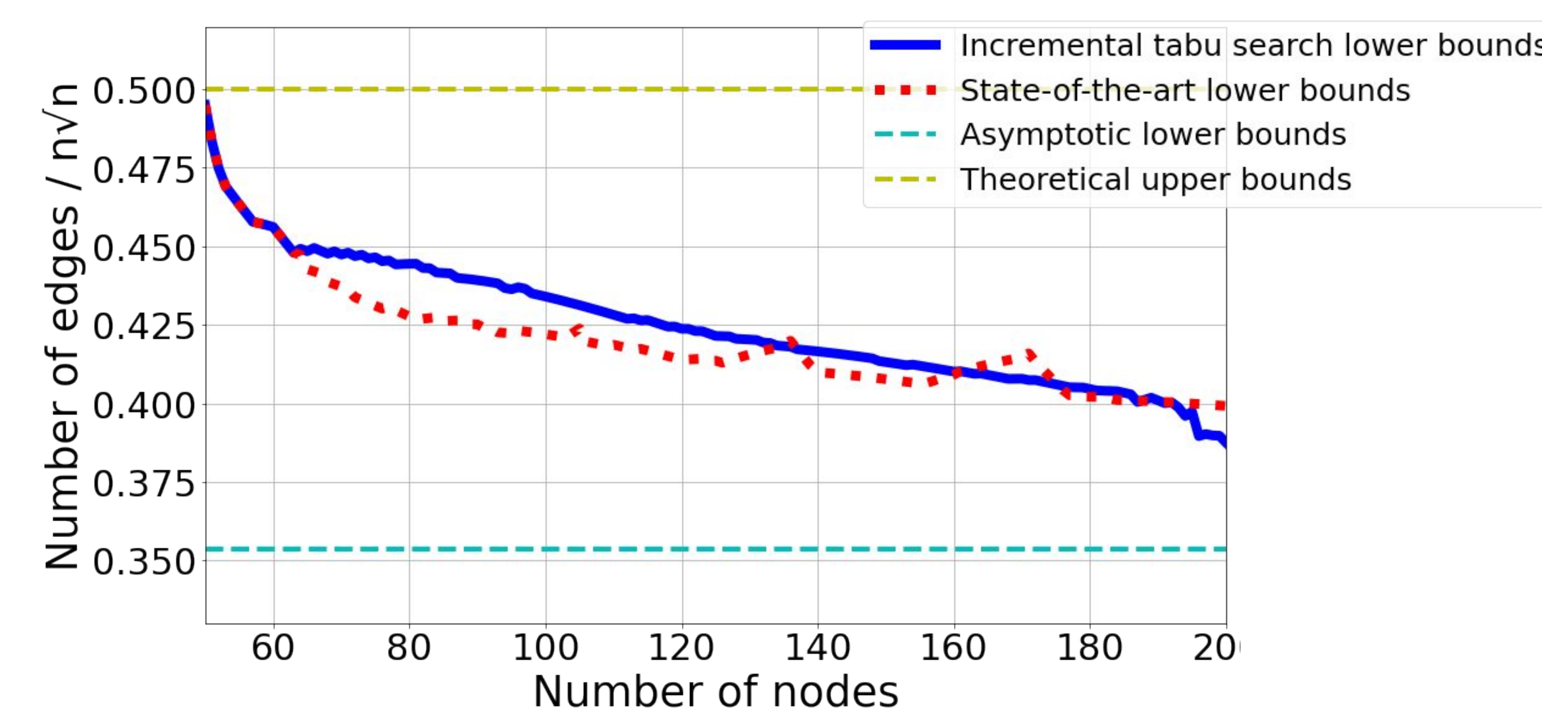

The gain from curriculum

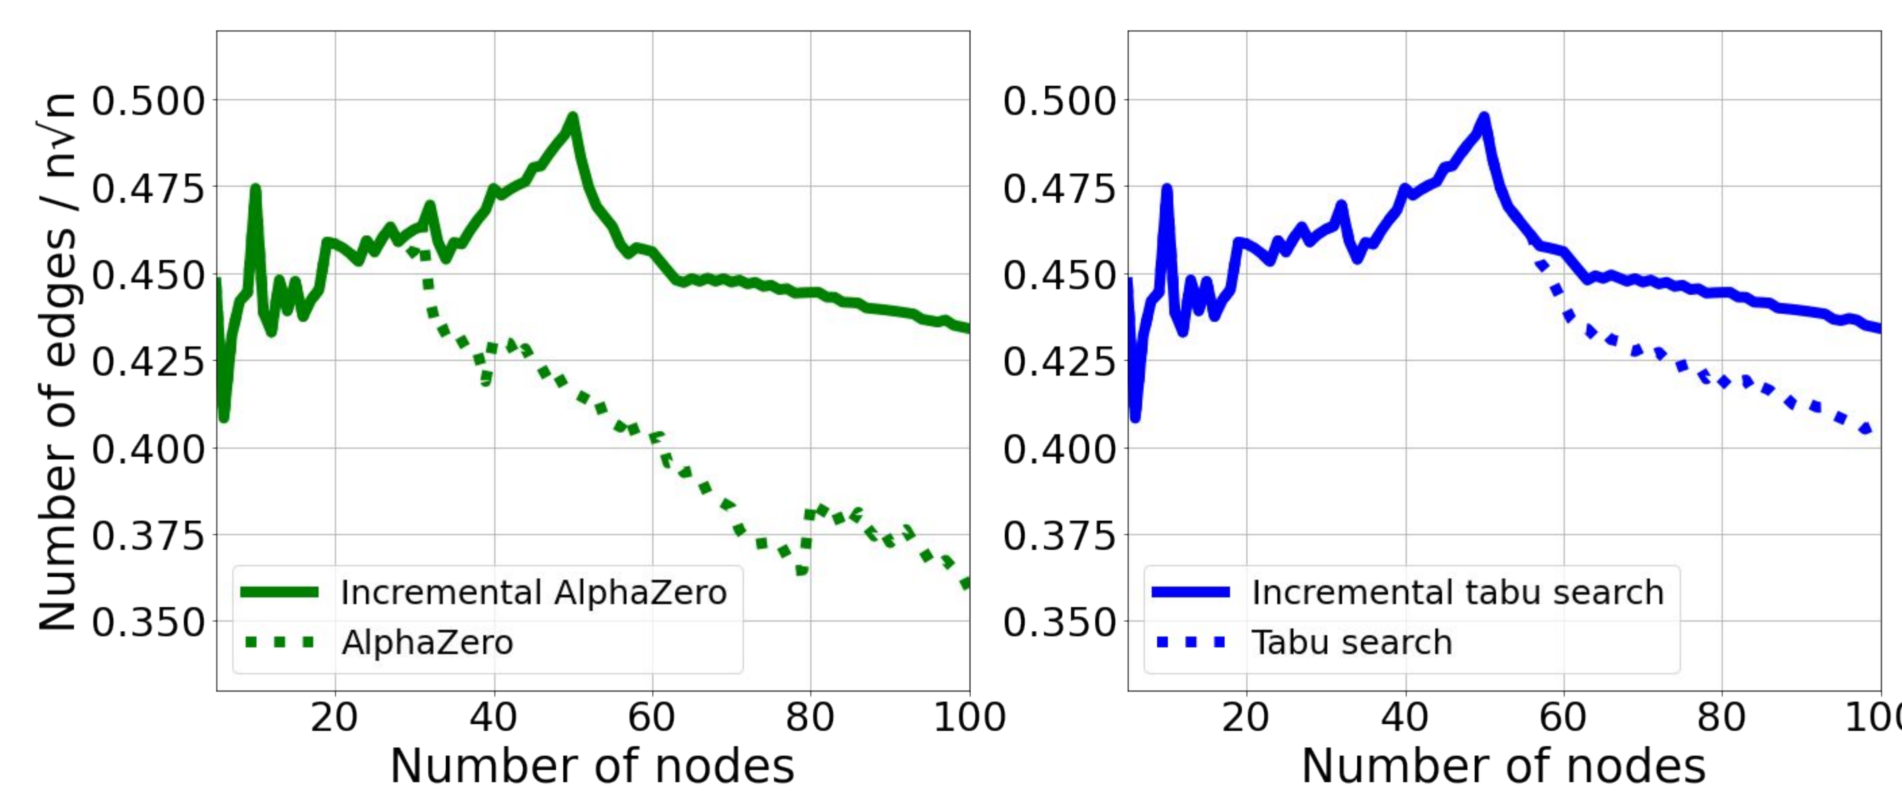

Contributions

- Proposed a new learning-to-search benchmark inspired by open conjecture in graph theory and obtained new lower bounds
- Compared AlphaZero with tabu search
- Showed that curriculum is key to this problem
- Designed a novel graph-generation environment and network architecture

Shortcomings & Directions

- Not much benefit of using learning over tabu search for improving bounds
- Tabu search explores much faster than network guided MCTS
- Not clear if maximizing expected return is useful on these problems
- How to combine tabu search with learning?