Reinforcement Learning for Placement Optimization

Anna Goldie and Azalia Mirhoseini

Google Brain
agoldie,azalia@google.com

ABSTRACT

In the past decade, computer systems and chips have played a key role in the success of artificial intelligence (AI). Our vision in Google Brain's Machine Learning for Systems team is to use AI to transform the way in which computer systems and chips are designed. Many core problems in systems and hardware design are combinatorial optimization or decision making tasks with state and action spaces that are orders of magnitude larger than that of standard AI benchmarks in robotics and games.

In this talk, we will describe some of our latest learning based approaches to tackling such large-scale optimization problems. We will discuss our work on a new domain-transferable reinforcement learning (RL) method for optimizing chip placement [1], a long pole in hardware design. Our approach is capable of learning from past experience and improving over time, resulting in more optimized placements on unseen chip blocks as the RL agent is exposed to a larger volume of data. Our objective is to minimize power, performance, and area. We show that, in under six hours, our method can generate placements that are superhuman or comparable on modern accelerator chips, whereas existing baselines require human experts in the loop and can take several weeks.

CCS Concepts/ACM Classifiers

- Computing methodologies → Machine learning
- Hardware → Electronic design automation

Author Keywords

Deep Learning, Reinforcement Learning, Placement Optimization, Device Placement, Reinforcement Learning for Combinatorial Optimization

BIOGRAPHY

Anna Goldie is a Staff Software Engineer at Google Brain and co-founder/tech-lead of the Machine Learning for Systems Team, which focuses on deep reinforcement learning approaches to problems in computer systems. She is also a PhD student in the Stanford NLP Group, where she is advised by Professor Chris Manning. At MIT, she earned a Masters of Computer Science, Bachelors of Computer Science, and Bachelors of Linguistics. She speaks fluent Mandarin, Japanese, and French, as well as conversational Spanish, Italian, German, and Korean. She has given high-profile keynotes in Mandarin Chinese, and her work has been covered in various media outlets, including MIT Technology Review and IEEE Spectrum.

Azalia Mirhoseini is a Staff Research Scientist at Google Brain. She is the co-founder/tech-lead of the Machine Learning for Systems Team at Brain, where they focus on deep reinforcement learning based approaches to solve problems in computer systems and meta-learning. She has a Ph.D. in Electrical and Computer Engineering from Rice University. She has received a number of awards, including the MIT Technology Review 35 under 35 award, the Best Ph.D. Thesis Award at Rice and a Gold Medal in the National Math Olympiad in Iran. Her work has been covered in various media outlets including MIT Technology Review and IEEE Spectrum.

REFERENCES

- [1] Anna Goldie and Azalia Mirhoseini, "Placement Optimization with Deep Reinforcement Learning", in *Proceedings of the International Symposium on Physical Design*, pp. 3-7, 2020.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

ISPD '21, March 22–24, 2021, Virtual Event, USA.

© 2021 Copyright is held by the owner/author(s).

ACM ISBN 978-1-4503-8300-4/21/03.

DOI: <https://doi.org/10.1145/3439706.3446883>