

Review 1

PC member: Marianne Ninet

Time: Nov 21, 19:35 UTC

Review

This extended abstract addresses an important question in Neural Architecture Search (NAS): whether analytical models can effectively replace direct hardware measurements for estimating latency and energy consumption during architecture optimization.

Main Strengths:

The problem statement is clear: measuring hardware metrics for every candidate model is expensive and time-consuming. The choice of ResNet18 on CIFAR-10 provides a manageable and well-understood testbed. Additionally, the introduction provides a good contextual overview on NAS.

Main Weaknesses:

First, while the method relies on "analytical performance models", the specific type is not mentioned. Similarly, clarifying the specific target hardware (or explaining why the specific choice doesn't matter) would strengthen the "hardware-aware" aspect mentioned in the abstract.

Second, the abstract mentions analyzing MnasNet (2019) and NASAIC (2020) to compare approaches, but these papers are barely discussed beyond the introduction. Explaining this comparison would be great.

Third, the distinction between "real data" and "simplified models" is not explicitly explained. Clarifying this would help the reader understand exactly which method is being utilized.

Regarding the experimental goals, the proposal would be stronger if it defined a clearer success metric. Finally, the search space in Table 1 is qualitative; providing specific numerical ranges rather than descriptions could improve the paper.

Quality of Writing

The writing is clear and well-structured. However, there is one minor typo in the abstract ("optimize" instead of *optimize*). Additionally, the term "effort-based learning" appears non-standard and could be defined.

Reviewer's confidence

2: Partly, I may be missing some concepts or elements of the state-of-the art, but I got the main idea

Usage of LLM

1: No, not at all

Confidential remarks for the program committee

(None provided)