Evaluation Report of Doctoral Thesis

Title: Heuristic Learning for Domain-independent Planning
PhD Candidate Name: RNDr. Ing. Otakar Trunda
Reviewer: Ing. Antonín Komenda, PhD. (Czech Technical University, Czech Republic)

The doctoral thesis of Otakar Trunda targets an interdisciplinary research topic of automated classical planning and deep machine learning. It is currently one of the most discussed interconnection of two areas of artificial intelligence from (without much exaggeration) the opposite ends of the spectrum. On one hand, classical planning fits into the symbolic artificial intelligence, stemming from the sequential decision-making grounds, providing highly efficient algorithms for providing solutions to (PSPACE-)hard combinatorial problems. The drawback is a requirement for precisely designed model of the problem. On the other hand, deep learning provides unparalleled efficiency in processing large amounts of data in problems where the model is unreachable or even non-existent. To allow for this, the methods are statistical and in contrast data-based, not model-based. Recently, a good deal of literature in the field of automated planning targets exactly this particular combination of research areas, as it is believed it could provide (semi-)automated gradual improvement of planners for problems closer to the real world. Therefore, the solved problem of the thesis is highly topical and important in broader context of artificial intelligence research.

The problem solved by the thesis aims at providing an automated learning process of heuristic functions for classical planning. This approach is well known in the field, however only recently tackled by deep learning and still without a generally well performing solution. The thesis clearly describes the solved problem and two groups of proposed solutions: (a) learned modifications of existing heuristics and (b) learned heuristic function using structural graph-based features detected in the planning problem definition. To my best knowledge, both approaches are novel, although the approach (a) with several important drawbacks (it is well-known that even small diversion from optimal heuristics can cause exponential increase in the states needed to be searched), which are however well explained in the thesis. The approach (b) is principally closest to the STRIPS-HGN by Shen, at al., but with focus not on particular action graphs, but the overall object graph. Such approach is the context of the solutions in literature, innovative and original. Said that, the particular principle of building of the object graph seems rather ad-hoc. The particular design decisions are motivated only in the form of restrictions (what and how the object graph could not be constructed) in contrast to a constructive form (why the object graph was designed in the way it is; ideally from a set of first principles targeting the properties of the learned heuristics and/or whether it describes some known planning structures, e.g., mutex groups, high-order facts, potentials or similar). The theoretical assurances are proven in form of bounds on the learned heuristic quality. Since the bounds target a general approach of learning a heuristic function, they are not particularly tight; however, they can be understood more as first steps in the direction of explainable and assurances on the learned heuristic, which is still an open problem.

As summarized in the previous paragraph, the results presented in the thesis are aiming at an important open problem, appropriately defined, with original, sound and formally well treated solutions. The experimental evaluation is deep enough and analyze the partial and final pieces of the designed solution from various angles and show well the limits of the proposed methods. Unfortunately, the final comparison of the learned heuristics is done only against the FF heuristic, which although still used for comparisons, can be hardly viewed as a fair comparison with the state of the art. At least the LAMA heuristic, or ideally red-black heuristics, should be included in the comparisons. More importantly, they could be used as additional information for the neural network to work with. Currently, the network is given only the FF estimates, but provided that it is fed with estimates from different heuristics, it could as a side effect provide an efficient learned weighted multi-heuristic/heuristic portfolio. However, in my opinion, the biggest problem of the thesis lies in the low-tier and low number of published papers by the student on the particular research topic. We are talking particularly about one CORE C conference paper (ICAART 2020) and it's post-proceedings (LNCS 2021).

The quality and formal treatment of the thesis are on good level. In general, I was missing a more thorough captions of several figures (e.g., I did not find explanation of the color coding in Figures 5.10, 5.11, 5.14, 5.22--5.26; Figure 5.7 is missing axes description). The text was well written and polished, the flow and structure of the thesis was good. There was only a small amount of vague formulations and/or imprecise statements. E.g., “the number of pair is too great”—how much great?, “Err(y, y)” vs. “Err(T, m)” seems like abuse of notation, “Since computing relaxed plans is still hard”—how much hard?.

Although the dissertation thesis is not backed by many relevant works published by the student, taking into the account the peculiarities of publishing interdisciplinary works, sound and clear message of the thesis, original solutions and detailed experimental evaluation, the work provides important step in the direction of utilization of deep learning for automated planning and therefore more general sequential decision-making by artificial intelligence.

Evaluation result: The dissertation work of Otakar Trunda shows the author **is able** of independent creative work.

In Prague, 2.2.2023

Antonín Komenda, Ph.D.