Deep Reinforcement Learning for Combinatorial Problems: A new approach for the 3DBPP container loading problem in logistics

J. Evers

HEC Liège - Management School of the University of Liège, QuantOM

e-mail: jevers@uliege.be

M. Schyns

HEC Liège - Management School of the University of Liège, QuantOM

e-mail: m.schyns@uliege.be

Keywords: Packing, Machine Learning, Artificial Intelligence

We will consider the case of the 3D Bin Packing Problem (3DBPP). It is classically formulated as follows: we have a set of containers where we need to store a set of cuboid boxes. Our objective is to place all the boxes in a minimum number of containers while avoiding exceeding the maximum capacity of each container and satisfying constraints of non-overlap between the different objects. This problem is NP-hard. Efficient heuristics are available for the 2DBPP, but the third dimension adds a significant level of complexity, at the combinatorial level itself, also with the addition of constraints that result from this third dimension in applications (management of fragility, stability, nature of the supports, weight distribution. . .).

While a traditional approach to tackling such problems is to resort to classical Operations Research techniques, some authors start considering techniques from other research streams, namely Artificial Intelligence (AI) and Machine Learning (ML). They have adapted and tested Reinforcement Learning (RL). In these approaches, an "agent" tests a possible sequence of actions and, in return, receives "rewards" from the "environment". By repeating sequences, the agent learns to recognize the most beneficial actions automatically. For complex problems, such as the one under consideration here, ML can be integrated into the framework. The approach is then renamed Deep Reinforcement Learning (DRL).

There is still much room for improvement. What can be observed from the literature on RL for the 3DBPP, is that a relatively small set of constraints has been considered so far. For example, besides the stability and orientation constraints that have already been incorporated in a few papers, many so-called "safety" and "logistic" constraints could be taken into consideration. We have also noticed that much of the literature on this topic is mostly focused on the single-bin configuration rather than the multi-bin setting.

Therefore, with this research, we plan on dealing with the 3 challenges outlined hereabove: the complexity of the problem which makes classical methods for the 3DBPP not efficient, the too-small number of constraints that have been addressed in recent literature on the subject, as well as the multi-bin setting which has been ignored up to now. To achieve this goal, we will build on an existing Deep Reinforcement Learning model for the 3DBPP we found in the literature, and determine a set of realistic constraints that could be added to this model while considering the multi-bin setting. This could also allow us to evaluate and compare the effectiveness of RL-based methods against well-known heuristics when a relevant set of constraints is considered.

Thus far, we have conducted a literature review and coded a first implementation of the model in Python. We hope that by the time of the conference, we will be able to present some results for a simple instance of the problem to add more constraints in the future.