Enhancing Autonomy in Warehouse Truck Unloading Via Greedy Optimization

Ashton Larkin
Brigham Young University

Fahad Islam, Anirudh Vemula, Maxim Likhachev
Carnegie Mellon University

Motivation

Unloading boxes from delivery trucks occurs frequently in warehouse settings. Manual truck unloading takes time and can physically injure employees. Automating this process with a robot can make truck unloading quicker and safer.

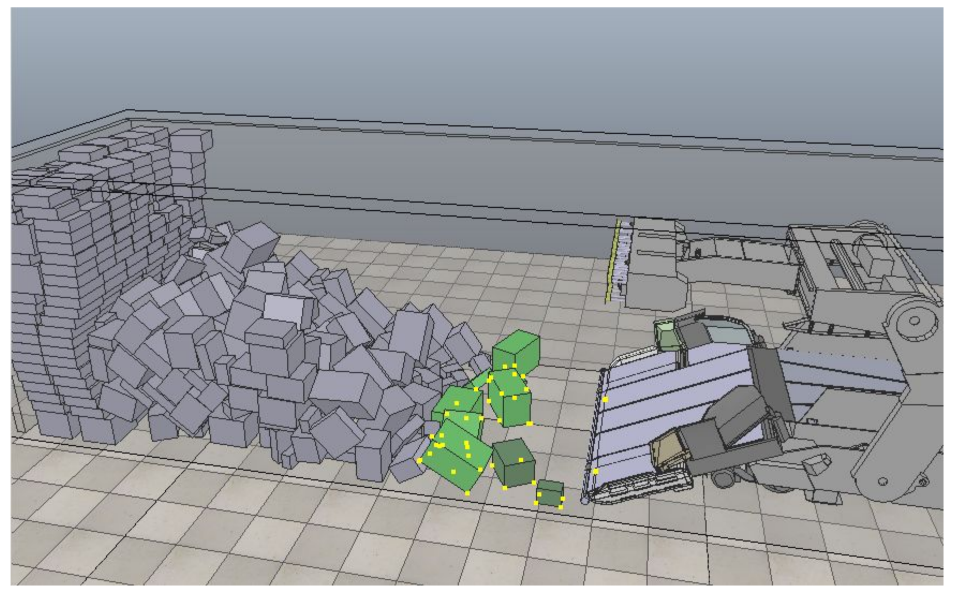

The simulation used for training and evaluation

Existing Approaches

- Planning under uncertainty
- Reinforcement learning

Technical Challenges

The physical interactions between boxes create a stochastic environment. Boxes vary greatly in size and mass, and box configurations in a truck depend on how a truck was packed.

Example environments the robot may encounter

The robot being used in this work [1]

Approach

Self-supervised learning using immediate reward optimization

Robot Actions:

- Pick left low, pick left mid, pick left high
- Pick right low, pick right mid, pick right high
- Sweep

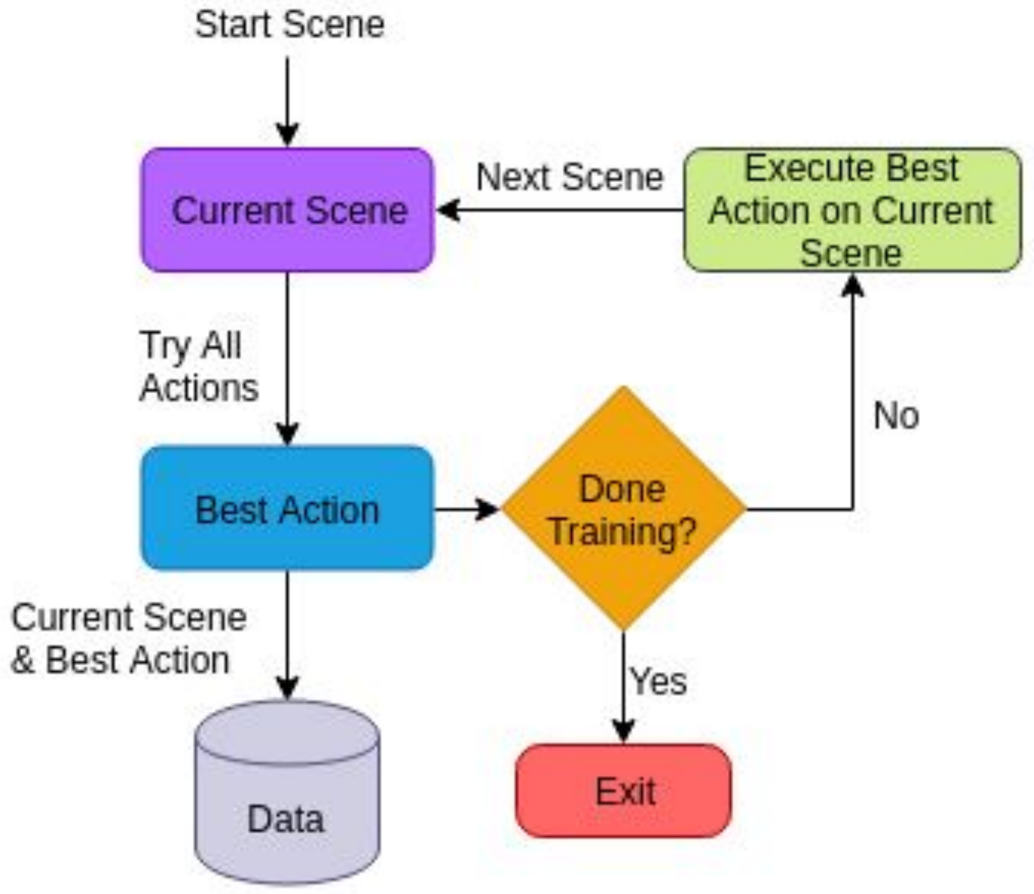

Training

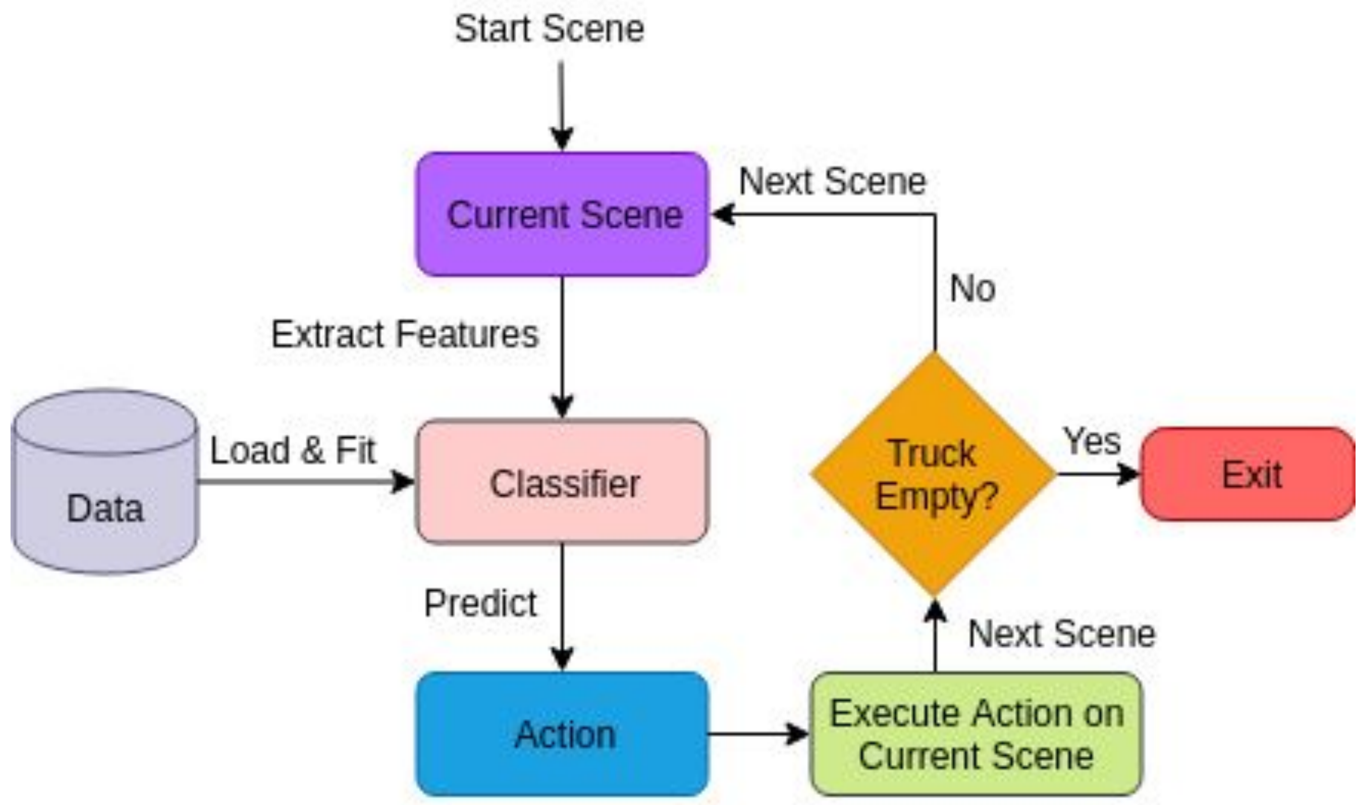

Evaluation

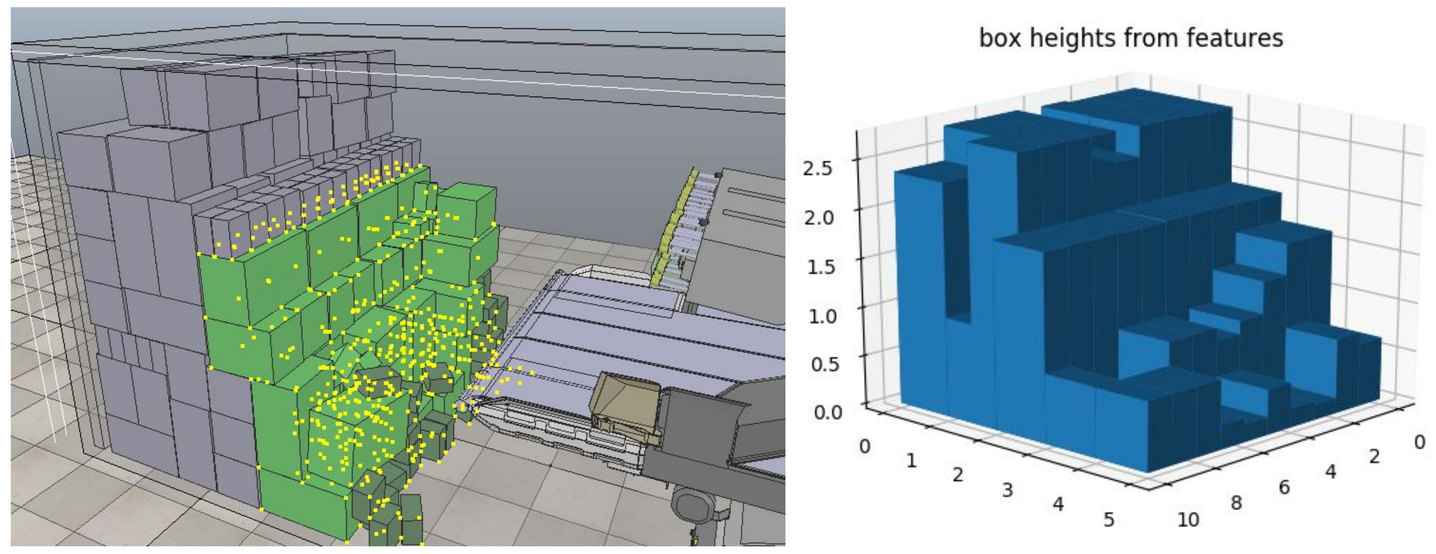

Feature extraction for a given scene

Experiments

4 environments used: A1, A2, B1, and B2.
A1 and B1 were used for training. A2 and B2 were used for evaluation.

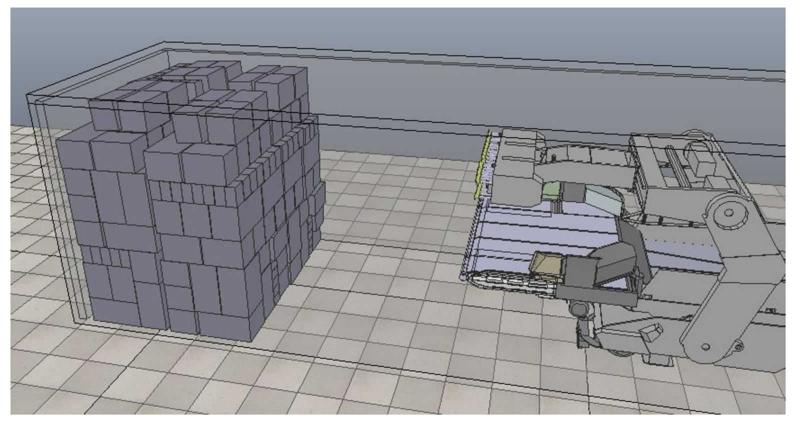

Environment A1

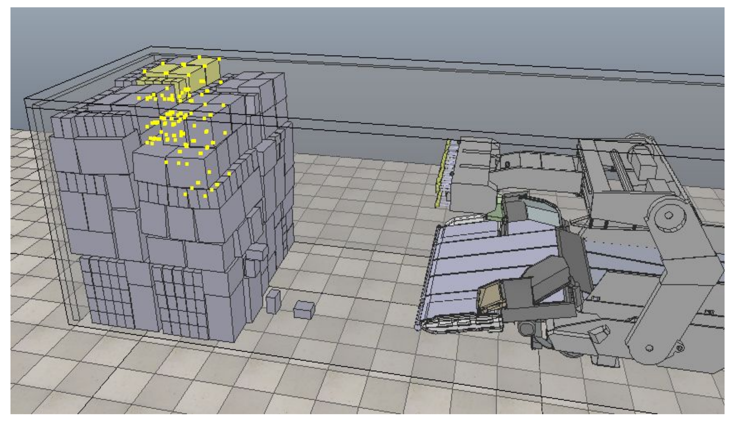

Environment A2

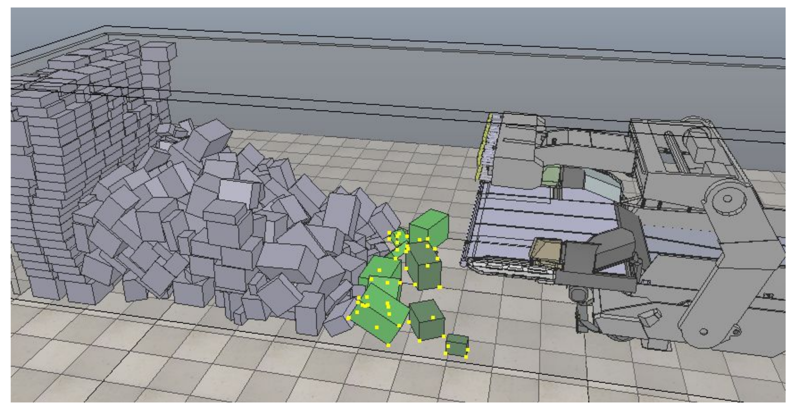

Environment B1

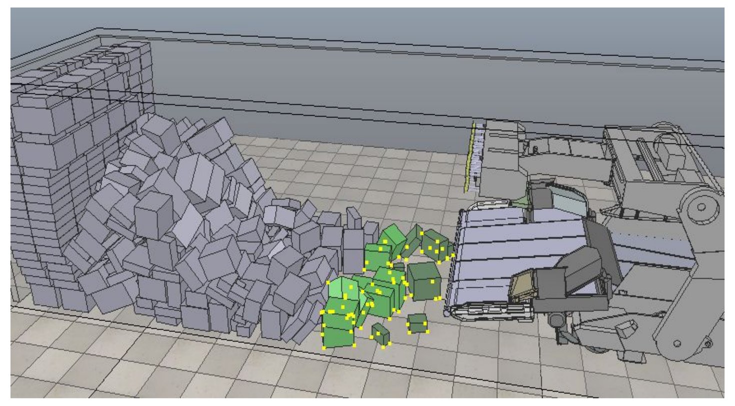

Environment B2

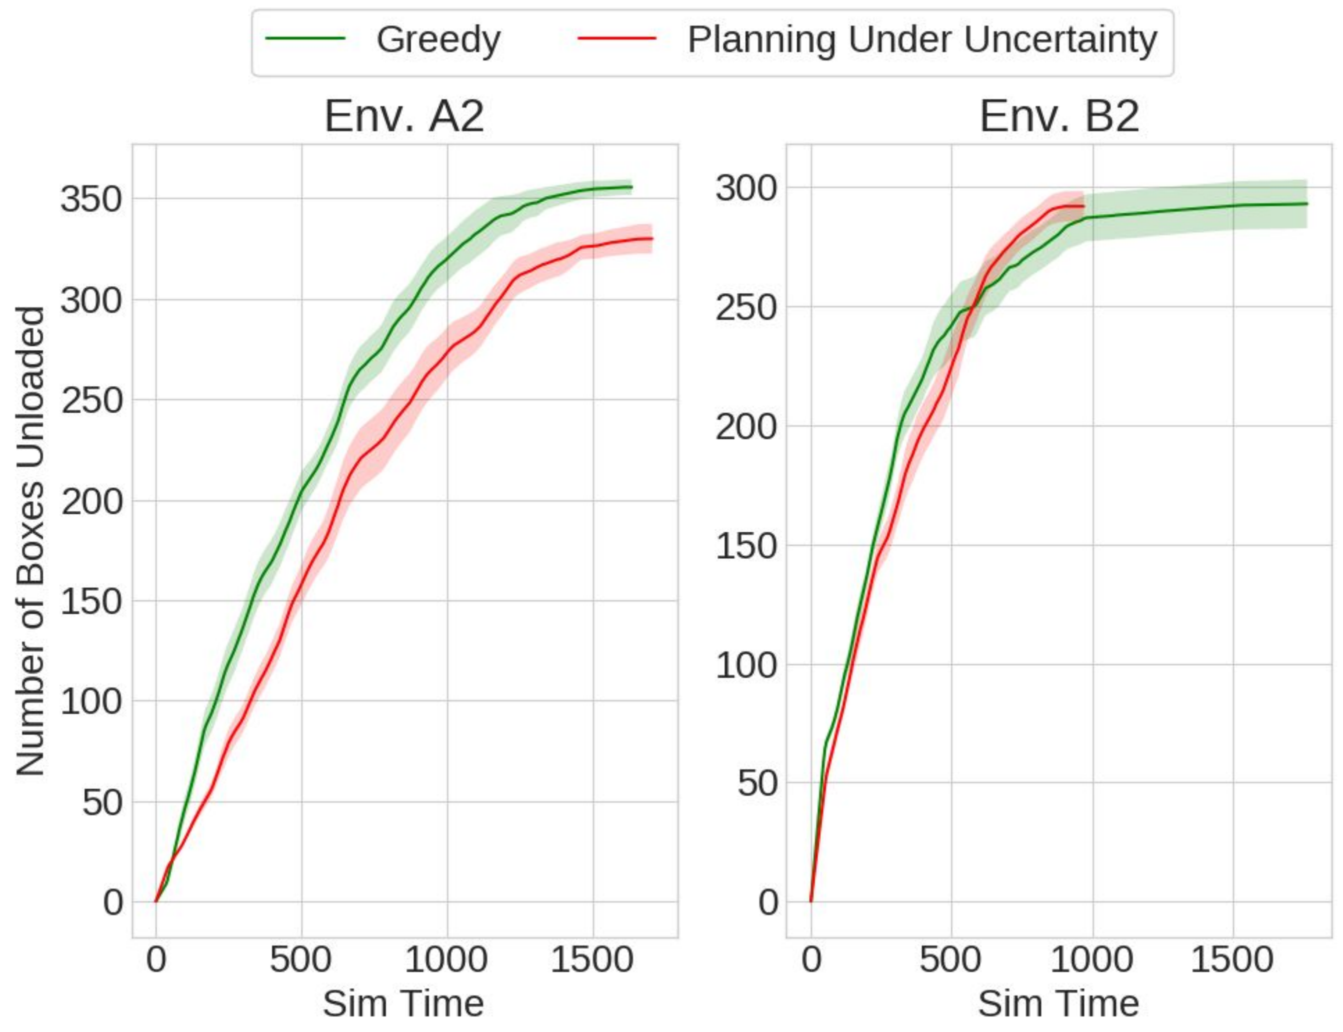

Our approach (green) compared to the previous approach (red)

Discussion

- Our approach outperforms a planning under uncertainty approach for structured environments
- Our approach performs similarly to a planning under uncertainty approach for unstructured environments
- Our approach is simpler and more computationally efficient than a planning under uncertainty approach
- Our approach works well since the problem space is bounded by the walls of the truck

Future Work

Multiclass cost-sensitive classification. In this work, we weighted the training data points based on the differences between the top two action rewards for a given scene. We would like to expand on this further by defining the classifier's loss as a function of the difference between the optimal action reward and predicted action reward.

Feature engineering. We used features similar to those used in the previous planning under uncertainty approach. These features do not seem to be informative for classification between task-level actions. Further work needs to be done in order to determine the proper feature representation for this problem.

Acknowledgements

Thank you to Carnegie Mellon University's Robotics Institute For Summer Scholars (RISS) program and the National Science Foundation for supporting this work. Ashton would like to give special thanks to Rachel Burcin, John Dolan, and the members of the Search Based Planning Laboratory at Carnegie Mellon University for their mentoring and guidance.

[1] https://www.bloomberg.com/news/articles/2019-05-03/robots-edge-closer-to-unloading-trucks-in-amazon-era-milestone