Multi-Robot Warehouse Optimization: Leveraging Machine Learning for Improved Performance

Mara Cairo
Bevin Eldaphonse
Payam Mousavi
Sahir
Sheikh Jubair
Matthew E. Taylor*
Alberta Machine Intelligence Institute
Edmonton, Canada

Graham Doerksen
Nikolai Kummer
Jordan Maretzki
Gurpreet Mohaar
Sean Murphy
Attabotics
Calgary, Canada

Johannes Günther†
Laura Petrich†
Talat Syed†
University of Alberta
Edmonton, Canada

KEYWORDS

Multi-Robot Systems; Supply Chain; Warehouse Management; Supervised Learning; Reinforcement Learning; Affinity Estimation

ACM Reference Format:

M. Cairo, G. Doerksen, B. Eldaphonse, J. Günther, N. Kummer, J. Maretzki, G. Mohaar, P. Mousavi, S. Murphy, L. Petrich, Sahir, J. Sheikh, T. Syed, and M. E. Taylor, Bevin Eldaphonse, Payam Mousavi, Sahir, Sheikh Jubair, Matthew E. Taylor*, Graham Doerksen, Nikolai Kummer, Jordan Maretzki, Gurpreet Mohaar, Sean Murphy, Johannes Günther†, Laura Petrich†, and Talat Syed†. 2023. Multi-Robot Warehouse Optimization: Leveraging Machine Learning for Improved Performance. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

1 INTRODUCTION

Supply chain issues, delays and shutdowns have dominated headlines, impacting individuals' ability to access, and companies' ability to deliver, crucial products and services. Changing consumer behaviours, accelerated by the Covid supply chain shock, have companies struggling more than ever to close the last-mile delivery gap. Attabotics Inc. [2] a Calgary-based robotics company that specializes in inventory management systems, offers a modern solution through its compact vertical warehouse structure and robotic order pickers. Attabotics replaces the rows and aisles of traditional fulfillment centers with a patented storage structure that uses both horizontal and vertical space, reducing a company's warehouse footprint by up to 85%. This empowers retailers, grocers, and e-commerce providers to place different sized fulfillment centers near high-density urban areas, decreasing carbon emissions by closing the last-mile delivery gap. With more than $165 million USD in investment, Attabotics's solution has been adopted by major brands, and has been featured in multiple venues like The Wall Street Journal, Time Magazine, and Tech Crunch.

* Corresponding author: matt.taylor@amii.ca
† Research originally performed while employed at Amii

We gratefully acknowledge funding from Scale-AI, as well colleagues who worked on earlier versions of this research, including A. Dankers, C. G. Jansson, and H. Nichols.

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

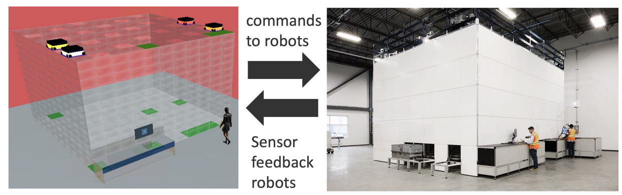

Figure 1: The system can control physical/virtual robots by directing them to bring bins to workstations, where physical/virtual humans or robots pick items to complete orders.

Attabotics recently engaged the Alberta Machine Intelligence Institute [1] (Amii) to leverage machine learning to help improve both speed and efficiency without sacrificing safety or reliability. The project focused on applying AI to make smarter decisions with the goal of increasing warehouse throughput by at least 25% to help grow profits and ensure customers get their products faster. Amii's expertise in machine learning has helped Attabotics improve throughput within these simulated warehouses using real-world data. We have also validated these performance improvements on a physical, in-house robot warehouse. Current work aims to move this proof of concept decision-making system to production.

Key Performance Indicator. While many metrics are possible, we choose to focus on *throughput*, the number of items delivered per time period. For a given set of orders, each of which can contain multiple counts of multiple SKUs (a Stock Keeping Unit is a unique ID assigned per product), how many SKUs per minute are packaged, on average? For any order set, maximizing throughput is closely related to minimizing total time needed to complete the orders.

Problem Definition. Robots are tasked to fulfill orders by bringing bins from the warehouse, to the workstations, swapping with the next bin, and charging when needed. Attabotics and Amii identified two main decision points where machine learning could be applied to increase throughput: (1) determining which SKUs should be placed in what bins and (2) deciding which bins robots should bring to workstations to fulfill orders as they occur.

2 APPROACH

Attabotics' control software runs in real time and can control both physical and simulated robots (i.e., a "digital twin" as shown in Figure 1). Developing and testing various solutions in software is

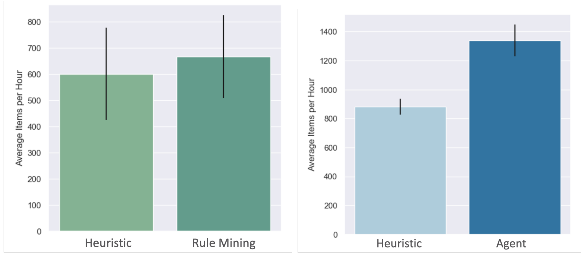

Figure 2: Results show that changes in throughput from using rule mining to assign SKUs to bins (left) and augmenting heuristic selection with data-driven inputs (right).

both faster and involves less wear-and-tear on the robots. A video (https://youtu.be/o_VHF6UVC5o) summarizes our approach.

2.1 Multi-Simulator Learning

The first insight was that rather than relying only on the true control software to provide data, which runs in real-time, we also developed a high-speed simulation environment. Modern machine learning techniques rely on data — the more data we have, the better. This new simulator (the *UltraSim*) was less realistic, but could run significantly faster, providing more data to help train our models. Furthermore, the fast simulation could be further sped up.

Consider a robot deciding what bin to select next. The control software could simulate all possible choices (and their consequences). Or, the fast simulator could more quickly simulate all the possible choices. If you could look infinitely far into the future, you could find the optimal decision, but this would take infinitely long in the real-world. This is exactly what supervised learning can be used for: given a certain decision about robot movement, rather than letting a simulation run and find out how long it would take the robot to finish its movement, a supervised learning model could predict how long it would take. Although imperfect, predictions allowed us to improve throughput, as discussed in Section 2.3.

The second insight was that the solution needed to be portable. We built a data pipeline so that our high-speed simulation, our bin allocation rules (Section 2.2), and supervised learning models (Section 2.3) could be quickly re-trained when robot capabilities improve, product demand or SKUs in the warehouse change, or a different warehouse layout is introduced.

2.2 Item and Bin Allocation

Exploratory data analysis highlighted several key factors that contribute to an increase in throughput. One of the most important was *bin compaction*, defined by the number of (relevant) SKUs delivered by a single robot for a given order. This section will briefly discuss how to put the right items into bins in the warehouse, while the next section will discuss which bins to pick for a given order.

Given historical data of processed orders, we applied the Apriori algorithm [3] to obtain association rules, which contain item-sets of various lengths. These rules can be ranked with their respective support values. Support of an association rule indicates the frequency of the item-set across all orders in the data. These rules are then cross-checked with physical constraints (e.g., 3D measurements and weight) of the items in the rules and the bins we intend to allocate them into. These filtered rules are then used to assign SKUs into bins in the warehouse, weighted by their support values.

2.3 Bin Selection and Order Fulfillment

Our decision-making agent combines machine learning with simple heuristics (e.g., number of lines completed) in order to make smarter choices about which bin to bring to a workstation. By using supervised learning over past robot travel in the warehouse, we can estimate the time needed to move a bin from any location to another. This allows us to decrease the total time it takes to fulfill an order by choosing bins that are both (1) quick to arrive and (2) contain relevant items. Note that this same supervised learning model is also used for state updates in the UltraSim.

3 RESULTS

Figure 2 shows the impact of bin compaction (left) and bin selection (right). Experiments are for two different customer warehouses with 7–10 days of historic data. Differences are statistically significant for bin selection ($p < 0.0001$ via paired t-tests) but not for bin compaction. Error bars show the throughput's standard deviation.

Performance of the newly developed systems will next be demonstrated at customer sites, moving from R&D to production. We expect the successes of robots in our internal test warehouse (and in simulation) will transfer to customers' physical warehouses.

4 CONCLUSION & FUTURE WORK

This work shows the potential of using machine learning to optimize space-efficient warehousing robotic solutions. By helping companies maintain space-efficient fulfillment centres closer to the customers they serve, and through innovative robotic solutions, Attabotics is helping to ensure the future stability in our local supply chains, helping them be more resilient to disruption.

Although successful, there are many avenues for future research.

- In addition to smart item-bin allocation, real-time or overnight bin replenishment will be required for long-term throughput performance.
- Seasonality and customers' buying pattern changes will require re-optimization. Showing continual optimization will further support our claim that this use of machine learning inside of robotic warehouses is game changing.
- We are currently developing an OpenAI Gym [4] environment where reinforcement learning [5] (RL) can be applied to make order fulfillment decisions with direct impacts in the real-world. We expect that an RL agent will be able to significantly outperform our data-driven heuristic approach.
- Instead of a single warehouse, we could consider extending to multiple warehouses of different sizes and locations as a much larger *network supply chain* problem.
- We have assumed a warehouse configuration is provided. The modeling techniques we have developed could also be used to help current or future customers "right size" new warehouses by understanding how different tradeoffs affect different KPIs.
- This project has focused on throughput. Due to the robustness of our data pipeline, we can always change the precise metric optimized for at the request of the customer (e.g., if different orders have different priorities).

REFERENCES

[1] 2022. Alberta Machine Intelligence Institute (Amii). https://www.amii.ca/about/. [Online; accessed 23-Nov-2022].
[2] 2022. Attabotics. https://www.attabotics.com/about. [Online; accessed 23-Nov-2022].
[3] Rakesh Agrawal, Heikki Mannila, Ramakrishnan Srikant, Hannu Toivonen, and A. Inkeri Verkamo. 1996. *Fast discovery of association rules.* AAAI Press, United States, 307–328.
[4] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. Openai gym. *arXiv preprint arXiv:1606.01540* (2016).
[5] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction.* MIT press.