LiveNet: Robust, Minimally Invasive Multi-Robot Control for Safe and Live Navigation in Constrained Environments

(Technical Report)

Will the Autonomous Mobile Robots Revolution Endanger Us?

(STS Paper)

A Thesis Prospectus
In STS 4500
Presented to
The Faculty of the
School of Engineering and Applied Science
University of Virginia
In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science in Computer Science

By
Srikar Gouru

November 8, 2024

Technical Team Members: Srikar Gouru & Siddarth Lakkoju

ADVISORS

Rohan Chandra, Department of Computer Science

Bryn E. Seabrook, Department of Engineering and Society

Prospectus

Introduction

Recently, the usage of autonomous mobile robots (AMRs) in social environments, such as warehouses, has grown immensely. It is estimated that 27% of warehouses currently use automated robotics, and that in the next five years, almost all warehouses will be automated (Law Offices of McMonagle Perri McHugh Mischak and Davis, 2023). However, these robots are by no means perfect. At Amazon fulfillment centers, serious injuries are 50% higher in centers that have mobile robots. In fact, in 2020 alone, there were 14,000 serious injuries caused by both stationary and mobile robots (BBC, 2020).

The issue lies in the fact that in these environments, AMRs have to be both safe and predictable. Safety ensures that the robots will not intentionally collide with humans or other agents in the environment. On the other hand, predictability allows humans to determine where the robots are going and adjust their own path to avoid them accordingly (Chandra et al., 2024). For example, when two people approach a door at the same time, one person usually slows down just enough that the other person can get through, and then the person that slowed down can follow, but no more than they need to. This type of motion is called *non-invasive* motion, since the robot is making the least amount of change possible to its original path to avoid the collision.

Existing methods to solve these issues rely on optimization-based controllers, using popular algorithms and techniques such as Model Predictive Control and Control-Barrier Functions (Chandra et al, 2024). However these optimization-based solutions are not deployable in the real world because they are computationally slow (Adebe, 2023). They are also often not robust enough to handle a variety of collision scenarios, and are instead usually hand-tuned to solve only a small set of scenarios. Additionally, Bogue (2024) discusses how the tasks that

AMRs are performing in warehouses, such as picking, sorting, and packing various items, are growing in complexity. This large complexity makes traditional AI algorithms difficult to use, and as more and more data is being gathered, much research is instead being focused on machine-learning based methods.

Handling these edge cases safely and running at real-time speeds are important when human lives are at stake. There currently exist standards of safety in warehouses for stationary machinery, including both hardware and software safeguards that must be in place for all technology (U.S. Department of Labor, 2024). Similarly, it is crucial for AMRs, which are often heavy, strong pieces of technology, to have similar software safeguards that follow safety protocols and maneuver predictability to avoid dangerous situations. A study by Bhattathiri et. al. (2024) discovered that participants preferred robots that exhibited human-like behavior, for example the usage of digital avatars, since they felt more comfortable and trustworthy. Similarly, if the AMRs exhibited human-like, *non-invasive* motion when avoiding collisions, it would help workers to adjust to these robots easier and feel safer. Autonomous vehicles need to communicate with people by behaving in an expected manner, similarly to how humans would behave (Hawkins, 2023). This technical project focuses on the design and implementation of a machine-learning based robot controller that exhibits safety, avoids collisions, and behaves *non-invasively* to help increase worker productivity and safety.

Technical Topic

Many existing warehouse systems use a centralized control system, in which a centralized computer dictates to each robot where it should move and what actions it should take.

However, these systems are prone to a single point of failure for the whole system and should only be used in extremely controlled environments (Rezaee & Abdollahi, 2018).

Existing decentralized methods to solve the issue of *non-invasiveness* rely on optimization-based controllers. These controllers work by mathematically formulating an objective, like trying to follow a path in a predictable manner, with certain mathematical constraints, like not colliding into humans and other robots, and try to solve this problem by guessing a solution and iteratively improving this guess (Murray, 2009). However, these more “traditional” control systems are subject to slow computation times, as the function to solve is extremely complex (Adebe, 2023).

Instead, I plan on using a machine-learning model to estimate this mathematical formulation that guarantees safety and predictability. Machine-learning models are significantly faster than traditional algorithms because the model can run many operations at once in parallel on a computer’s graphics processing unit (GPU), allowing it to run much faster than the iterative algorithms that optimization-based controllers use (Yi, 2024). Additionally, unlike optimization-based control systems, they can be trained on various scenarios and can internally learn and extrapolate on how to handle edge-cases or unexpected environments (Adebe, 2023).

A major issue with machine-learning systems in robotic controls is the lack of guarantee in safety due to the network being a black box with no mathematical guarantees. To overcome this hurdle, I intend on utilizing the recently released BarrierNet, a machine-learning framework that is capable of learning certain safety functions, such as avoiding collisions with other vehicles (Xiao et al., 2021). This self-learned safety function, along with the model being trained on safe data, will allow the network to extrapolate the data and behave safely even in scenarios that it was not trained in.

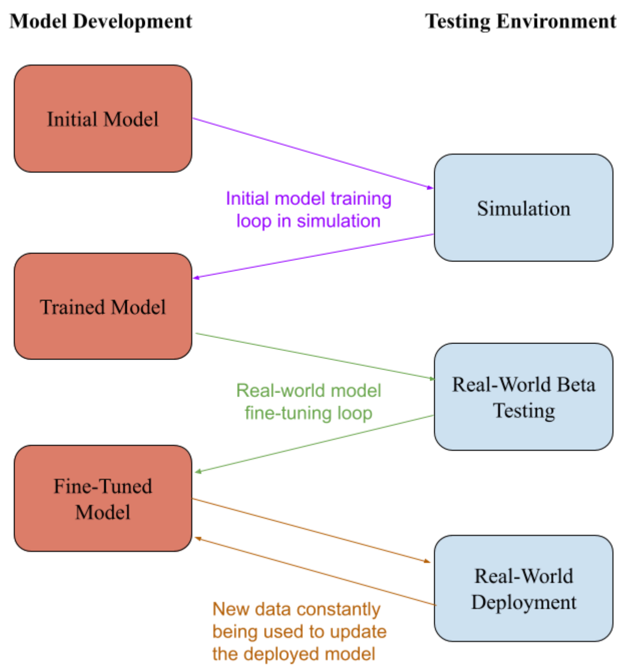

Figure 1: The training, testing, and deployment cycle of a machine-learning robotics model.

The safety and predictability of this new technology would be ensured through iterative testing of the software, first in simulation, and then in real-world scenarios (see Figure 1). Additionally, using software alerts can help keep track of certain metrics like safety score and time-to-collision internally. The use of machine-learning models allows the algorithm to be tuned to minimize these metrics. Another advantage of using a machine-learning based approach is that it scales well, since as more robots are deployed, more data gets generated that the rest of the robots can be trained on (Kober, Bagnell, & Peters, 2013). Any novel scenario that a single robot runs into will be added to the shared dataset that every other robot can train on and learn from.

STS Topic

AI-enhanced autonomous mobile robots (AMRs) are being rapidly adopted in industrial settings to improve productivity and overcome various technical challenges (Bogue, 2024). However, these environments have historically been designed for human use, and the integration of these complex systems without careful consideration could result in dangerous and unsafe environments. Using the framework described by Susan Leigh Star in her 1999 article *The Ethnography of Infrastructure*, we can analyze the integration of these AMRs into human-based industrial settings.

A core aspect of Star's Ethnography of Infrastructure Framework is that new technology is often built on an existing base. The design of many warehouses and similar industrial areas are optimized for human logistics, including the width of the walkaways, the size of the building, and the weight of the boxes that need to be transported. The AMRs that are developed should be designed specifically for a human-based warehouse, rather than considering humans as an extra task. It is important that these robots interact with people and other robots in the fulfillment centers and make safe, logical decisions when doing so.

Another crucial aspect of the Ethnography of Infrastructure Framework is that new technology that is integrated should follow the existing standards in place. For AMRs, this involves the ability to act both safely and reliably in social situations when interacting with humans. However, machine learning is often avoided in robotic control systems due to its lack of transparency as they are often seen as black boxes. Many studies are being conducted in attempts to develop more interpretable or explainable machine-learning models, however this concept of *explainable machine-learning* is still an open question and has not been fully researched

(Doshi-Velez & Kim, 2017). On the other hand, many complex problems cannot be solved efficiently using traditional algorithms. Instead, the speed, robustness, and something can often outweigh the lack of explainability and lead to a more reliable system than traditional algorithms (Jordan & Mitchell, 2015). Designing a system that balances trustworthiness and explainability with performance and reliability is crucial to engineer these algorithms for safety-critical systems.

Additionally, industrial settings have standards not only for safety, but also for privacy. According to the UK's Information Commissioner Office (ICO), the access to CCTV footage in warehouses is limited to a few, trustworthy people. Additionally, policies like the Data Processing Agreement (DPA) limit data storage length, requiring the National ANPR Data Center to perform 6.7 billion data deletions annually (Bamford, 2015). These limitations on data access and retention clash with machine-learning concepts of open source datasets, which are used to share data across the world and provide external developers the opportunity to help improve these machine learning models. A data pipeline must be developed that allows the technology to improve without violating worker privacy.

Another key aspect of infrastructure that Star mentions is being noticeable when something isn't working as intended. When integrating with existing, high-risk environments surrounded by people, it is important to have built-in fallbacks that ensure safety. Automated data logging and alert pipelines can be used to ensure that the system is functioning as intended (Fergus, 2023). These systems will also increase the trust that workers have in the robots, as they can visualize externally that the robots are working as intended.

Research Question and Methods

The inherent safety concerns with AMRs and their interaction with humans in warehouse spaces poses an important question about whether or not their utility outweighs their potential harm. That is, how does the increase in collaborative robots in industrial settings affect workers' productivity and safety in these environments?

To investigate this issue, I plan on analyzing historical data regarding injury rates and worker output per-capita before and after the integration of AMRs to get metrics on safety and productivity improvement. Using reports published by the Occupational Safety and Health Administration (OSHA) and the North American Industry Classification System (NAICS), I can empirically measure these metrics. These metrics will help demonstrate the safety and efficiency standards before the integration of AMRs, , which is a key factor in any infrastructure integration (Star, 1999).

In addition to organizational reports, I plan on interviewing non-robot (people) workers who currently interact with these robots in the industry. This includes a variety of jobs that use AMRs for both such as warehouse associates who use AMRs to carry heavy objects, healthcare staff who use AMRs to decrease pathogen exposure (Sahoo & Choudhury, 2023), and restaurants that serve food using robot waiters to enhance customer experience. I intend to ask these workers, anonymously, about their experience with the robots in terms of work efficiency and task completion. I also plan on asking if any injuries have been reported since the robots were introduced into their workplace. Gathering data from workers from industries that use AMRs will supplement the quantitative data that the organizational reports provide with a qualitative, holistic view of the workers' views on the technology.

Conclusion

The increasing use of AMRs in warehouses raises significant concerns for worker safety and productivity. These robots have the potential to significantly improve productivity, but must behave safely and predictably to ensure worker well-being. The technical deliverable is a safe and non-invasive control system for AMRs that utilizes machine learning to navigate in socially complex environments. The STS deliverable of the research conducted in this project are insights on how to safely and reliably integrate novel AMR technology into existing industrial environments originally designed for humans. The expected results of this research will help shed light into the impact of AMRs on workplace safety and productivity, particularly through worker surveys and historical data analysis. This project will help make AMR technology safer and more reliable for usage in human-based environments.

References

Adebe, S. (2023, February 10). *A comparison of model-based MPC and Machine Learning Mpc*. Medium. https://medium.com/@airob/a-comparison-of-model-based-mpc-and-machine-learning-mpc-9e051d689008

Bamford, J. (2015). (rep.). *A data protection code of practice for surveillance cameras and personal information*. Information Commisioner's Office. Retrieved November 8, 2024, from https://ico.org.uk/media/1043340/surveillance-by-consent-cctv-code-update-2015-jonathan-bamford-20150127.pdf.

BBC. (2020, September 30). *Amazon Warehouse Robots "increase staff injuries."* BBC News. https://www.bbc.com/news/technology-54355803

Bhattathiri, S. S., Bogovik, A., Abdollahi, M., Hochgraf, C., Kuhl, M. E., Ganguly, A., Kwasinski, A., & Rashedi, E. (2024). Unlocking human-robot synergy: The Power of Intent Communication in warehouse robotics. Applied Ergonomics, 117, 104248. https://doi.org/10.1016/j.apergo.2024.104248

Bogue, R. (2024), "The role of robots in logistics", Industrial Robot, Vol. 51 No. 3, pp. 381-386. https://doi-org.proxy1.library.virginia.edu/10.1108/IR-01-2024-0022

Chandra, R., Zinage, V., Bakolas, E., Stone, P., & Biswas, J. (2024). *Deadlock-Free, Safe, and Decentralized Multi-Robot Navigation in Social Mini-Games via Discrete-Time Control Barrier Functions*. https://doi.org/10.21203/rs.3.rs-3979309/v1

Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning.
ArXiv preprint arXiv:1702.08608.

Fergus, S. (2023). A complete guide to automated data pipelines. Retrieved from https://www.shipyardapp.com/blog/automated-data-pipelines/

Hans. (2024, January 10). *The Future of Warehouse Technology: Preparing Employees for amrs*. MARC Mobile Autonomous Robotic Cart. https://www.multechnologies.com/blog/preparing-employee-for-amrs

Hawkins, A. J. (2023, October 13). *How will driverless cars "talk" to pedestrians? Waymo has a few ideas*. The Verge. https://www.theverge.com/2023/10/13/23913251/waymo-roof-dome-communicate-intent-pedestrian-driver

Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, *349*(6245), 255-260.

Kober, J., Bagnell, J. A., & Peters, J. (2013). Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, *32*(11), 1238-1274.

Law Offices of McMonagle Perri McHugh Mischak and Davis. (2023, January 16). *Does robot automation mean more warehouse worker injuries?*. McMonagle Perri McHugh Mischak and Davis. https://www.mpmpc.com/blog/does-robot-automation-mean-more-warehouse-worker-injuries

Murray, R. M. (2009). *Optimization-Based Control*. *Control and Dynamical Systems*. Retrieved from https://www.cds.caltech.edu/~murray/books/AM05/pdf/obc09-obc09_03Mar09.pdf

Sahoo, S. K., & Choudhury, B. B. (2023). Challenges and opportunities for enhanced patient care with mobile robots in Healthcare. *Journal of Mechatronics and Artificial Intelligence in Engineering*, *4*(2), 83–103. https://doi.org/10.21595/jmai.2023.23410

Star, S. L. (1999). The ethnography of Infrastructure. *American Behavioral Scientist*, *43*(3), 377–391. https://doi.org/10.1177/00027649921955326

Rezaee, H., & Abdollahi, F. (2018). Decentralized control in robotic swarms. *Swarm Intelligence - Volume 1: Principles, Current Algorithms and Methods*, 475–507. https://doi.org/10.1049/pbce119f_ch17

U.S. Department of Labor. (2024). *Robotics - standards*. Occupational Safety and Health Administration. https://www.osha.gov/robotics/standards

Xiao, W., Hasani, R., Li, X., & Rus, D. (2021). Barriernet: A safety-guaranteed layer for neural networks. *arXiv preprint arXiv:2111.11277*. https://doi.org/10.48550/arXiv.2111.11277

Yi, X. (2024). A Study of Performance Programming of CPU, GPU accelerated Computers and SIMD Architecture. *arXiv preprint arXiv:2409.10661*.