DARKO-Nav: Hierarchical Risk- and Context-aware Robot Navigation in Complex Intralogistic Environments

Elena Stracca[1*], Andrey Rudenko[2*], Luigi Palmieri[2*], Paolo Salaris[1], Luca Castri[3], Nicolò Mazzi[4], Vasilije Rakcevic[5], Narunas Vaskevicius[2], Timm Linder[2], Nicola Bellotto[3,6], Tim Schreiter[5], Yufei Zhu[7], Manuel Castellano Quero[7], Olga Napolitano[1], Elisa Stefanini[1], Lukas Heuer[2], Martin Magnusson[7], Abdalla Swikir[5], and Achim J. Lilienthal[5,7]

[1] Research Center "E. Piaggio", dept. of Information Engineering, University of Pisa, Italy, `elena.stracca@phd.unipi.it`
[2] Corporate Sector Research and Advance Engineering, Robert Bosch GmbH, Germany, {`andrey.rudenko, luigi.palmieri`}`@de.bosch.com`
[3] University of Lincoln, UK
[4] Spindox Labs, Italy
[5] Technical University of Munich, Germany
[6] University of Padua, Italy
[7] Centre for Applied Autonomous Sensor Systems, Örebro University, Sweden
* Shared first authorship

Abstract. We propose a flexible hierarchical navigation stack for a mobile robot in complex dynamic environments. Our solution includes several layers to efficiently decompose the complex problem into task planning, navigation, control and safe velocity components. In contrast to the prior art, our system at every level incorporates diverse contextual information about the environment, anticipates navigation risks and proactively avoids collisions with dynamic agents.

Keywords: hierarchical navigation, intralogistic robotics

1 Introduction

In the last years there has been a large increase of mobile robots working in both industrial and service sectors. In such settings, mobile robots will need to enter and quickly integrate into the existing facilities, navigate in inherently human environments, interact with customers and human co-workers, and gradually assume increasingly more complex tasks involving efficient delivery and mobile manipulation. This vision creates unique challenges to the navigation system, which needs to drive the robot in large, complex and dynamic environments, while satisfying strong safety requirements. The challenge goes beyond the isolated problems of avoiding collisions with static obstacles or moving through a crowd, and requires a systematic approach to reduce the planning complexity and incorporate the diverse contextual cues with different time scales and representations. In shared dynamic environments, layered navigation frameworks provide an efficient approach to address complex navigation tasks [15]. A safety architecture to enable human-aware navigation in industrial environments, integrating multiple safety layers, was developed in [8].

In this work, we present a novel hierarchical navigation architecture concept designed to improve both performance and safety by incorporating explicit reasoning around contextual cues, causal inference, heterogeneous risk factors, and learned motion patterns of moving agents across multiple layers of the navigation pipeline. Our solution combines a policy and route decision-making module that minimizes task failure and performance degradation risks in the presence of humans, a global planner to estimate the collision risks coming from the environment dynamics patterns, a predictive context-aware MPC planner to drive robot in dynamic and cluttered environments, and a safe motion unit to control velocity in immediate proximity to people. This work has been developed as part of the EU project DARKO[1]. Most system components have been successfully integrated and tested in several qualitative experiments on the DARKO platform (a Robotnik Summit-XL steel mobile base equipped with a manipulator).

2 Method

2.1 Requirements

We have developed our architecture according to the following requirements:

- Make use of the many *contextual cues* that the robot can perceive. These cues include detected positions, poses and activities of the moving people, static and dynamic scene maps, congestion risks estimations and human-robot spatial interaction components.
- Decompose the complex planning problem into a *hierarchical system* of layers, each effectively handling a different planning horizon from risk-aware task scheduling to safe velocity control in close proximity to people.
- Make the navigation safe and efficient by considering *predictions and risk assessments* at every layer. The generated plans should reduce the risk of collision or in general robot unsafe operation to a minimum.

According to these requirements, we have developed a multi-layered navigation stack, see Fig. 1. The architecture follows a *predictive planning* setup: differently from a traditional *sense-plan-act* one [13], our architecture plug-ins various types of predictions of surrounding humans and dynamic objects inside the layers, covering several time scales. Similarly we allow the usage of several contextual cues across the several layers, e.g. collision risks and human activities. Our architecture is composed of four main layers: (1) Reasoning and Scheduling, (2) Global planning, (3) Local planning, (4) Safety layer for vehicle motion.

[1] https://darko-project.eu/

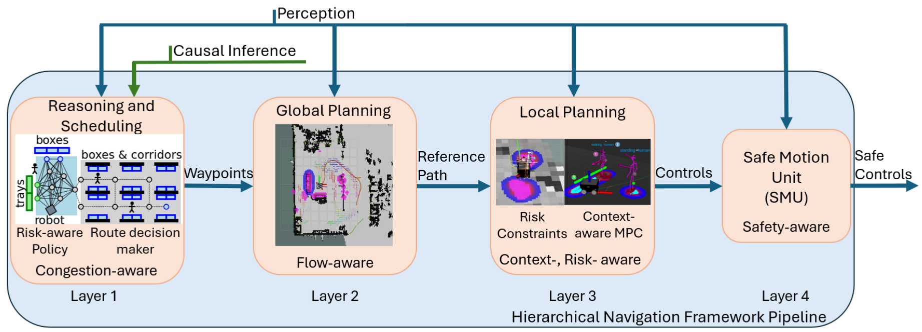

Fig. 1: Architecture overview of our navigation system

2.2 Components

Perception DARKO-Nav relies on a robust and efficient 3D human perception system to provide contextual cues for its hierarchical navigation layers. The DARKO perception stack utilizes a diverse set of on-board sensors such as 3D lidar, RGB-D and fisheye cameras to provide 3D detection and tracking, skeleton pose estimation and activity recognition. For real-time 3D human centroid detection, we utilize a TensorRT-accelerated variant of RGB-D YOLO [7]. The resulting detections are then processed by a top-down absolute 3D human pose estimator based on volumetric heatmaps [12]. Human trajectories are tracked using a Kalman filter-based tracker [6] and together with temporally associated 3D skeletons are subsequently fed into a human activity classifier, implemented with an RBF kernel SVM trained on skeletal features. Predicted activity classes are integrated over time through fixed-lag, mode-based filtering.

Causal Inference In addition to the individual human detection and tracking, DARKO robot is equipped with a causal inference-based module to model the spatial behaviors of people and the robot. These causal models consider various contextual factors, such as time, human-robot proximity, people density, and the robot's tasks. The module employs ROS-Causal [2] to collect human and robot trajectory data and conduct onboard reasoning. Causal discovery is performed by F-PCMCI [3] and CAnDOIT [4], which are specifically tailored for robotics applications. The ultimate goal is to create a compact causal representation of the environment that can be used for prediction and estimation by the risk analysis and intention communication components of DARKO-Nav.

Layer 1: Reasoning and Scheduling The task of the top layer is to define the next operational actions, such as selecting which object to retrieve or engage with and determining strategic waypoints for efficient task execution. The aim is to establish high-level movement and scheduling strategies that guide the robot efficiently through the environment, optimizing task completion and minimizing risks (e.g., navigating through congested areas). To do so, we introduce a discrete set of waypoints where the robot can perform (manipulation) actions to complete its mission. The goal is to obtain a risk-aware policy that guides the robot through a sequence of (potentially fallible) waypoint navigation and manipulation tasks. This policy is obtained via an Approximate Stochastic Dynamic Programming framework, which inherently models potential task failures and their consequences. Additionally, we implement a risk-aware optimal routing module for extended navigation tasks, providing a policy to select routes that minimize performance risks, such as delays from necessary slowdowns in human-robot interactions. Formulated as a Markov Decision Process, this problem models dynamic replanning, obstacle observability, alternative route costs, and encounters severity, creating a more proactive navigation strategy. For long-term task planning and risk-aware routing, we leverage the previously described causal inference module. The model derived from causal inference predicts the density of people at the various waypoints within a specified time horizon, allowing the modules to make risk-informed decisions.

Layer 2: Global planning Once the task of the robot is defined and the sequence of waypoints is available, this layer is responsible for guiding the robot between them considering the long-term human behaviors and dynamic flows in the environment. In that regard, we represent generalized human flows using Maps of Dynamics (MoDs [5]), which encode the direction and intensity of human motion in each location, allowing the robot to naturally fit into the flow or avoid the areas that are likely to be crowded and turbulent. To that end, we use MoD-aware planning, e.g. CLiFF-RRT* [11] to generate a path by exploring the configuration space via random sampling of robot configurations, using the MoD-encoded human movement directions and intensities as sampling priors.

Layer 3: Local planning To enable safe and efficient navigation in shared, dynamic environments along the global path, we propose a local planner based on a context- and risk-aware Model Predictive Control (MPC). Our formulation [14] incorporates diverse environmental and human-related information, such as full-body 3D human skeleton poses projections, human activity detection, and 2D motion predictions, allowing the robot to proactively avoid walking humans and bypass stationary ones. To integrate the notion of risk in this layer, the MPC leverages heterogeneous risk maps, comprising a collision risk map that accounts for dynamic uncertainties in obstacle detection and localization, and static risk maps like alignability maps [1], which encode the risk of localization accuracy loss for all the traversable regions. Additionally, an active sensing module [9] enhances safety by reducing collision and localization uncertainties through optimized information flow [10], enabling dynamic re-planning to mitigate task failures and aid in recovery from localization loss. The MPC incorporates risk maps as soft constraints with adjustable weights based on risk severity, and we plan to include the active sensing module similarly.

Layer 4: Vehicle Safe Motion Unit Initially developed for ensuring human safety during pHRI with stationary robot arms, the Safe Motion Unit (SMU) approach limits the robot velocity according to an injury data-driven safety framework. Given relative human-robot distance and velocity information and

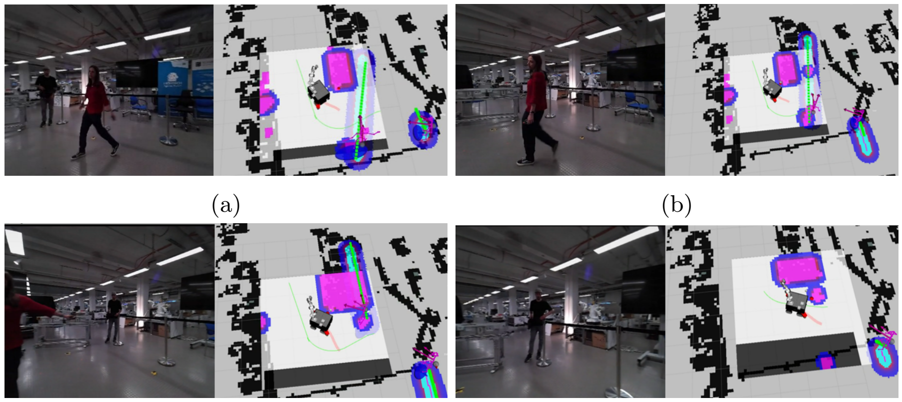

(a) (b)

(c) (d)

Fig. 2: RViz visualization and onboard camera view of the robot in a shared workspace. Human skeleton projections are shown as blue ellipsoids, with predicted trajectories indicated by green dotted lines with associated uncertainties. The grey map shows the static risk area, while the pink and blue map is the collision risk map. The path generated by the context- and risk-aware MPC is shown in red.

based on a rich injury database for various impact curvatures, the vSMU shapes the robot motion such that the impact energy of possible dynamic collisions cannot exceed the safety limits set by the standard or other safety mappings chosen to be respected.

3 Evaluation, Conclusions and Future Work

This paper outlines our concept of the navigation system for the DARKO robot, which is expected to safely traverse long distances in topologically-complex dynamic environments. Preliminary evaluation of the individual components and their combinations, except for the route decision maker and the active sensing module, has shown promising results in real-world experiments. Fig. 2 shows the robot moving during integration tests at KI Fabrik in Munich. The robot moves taking into account human 3D poses, activities, trajectory prediction and risk maps. As shown in Fig. 1, the robot is able to plan the global path according to the typical human motion direction in this part of the environment. In another real-world experiment, aimed at testing the localization risk-aware capabilities of the local planer, the robot managed to successfully avoid high-risk areas, potentially leading to large localization errors, thanks to the integration of an alignability map of the environment. Towards the final milestone of the DARKO project in June 2025, we plan to conduct extensive testing in dynamic environments with increased complexity, focusing on metrics such as task efficiency, adherence to safety specifications, and adaptability to human behavior.

Acknowledgments: This work was supported by the European Union's Horizon 2020 research and innovation program under agreement 101017274 (DARKO).

References

1. Manuel Castellano Quero, Tomasz P. Kucner, and Martin Magnusson, *Alignability maps for the prediction and mitigation of localization error*, IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS) Workshop on Closing the Loop on Localization, 2023.
2. Luca Castri, Gloria Beraldo, Sariah Mghames, Marc Hanheide, and Nicola Bellotto, *Experimental evaluation of ros-causal in real-world human-robot spatial interaction scenarios*, IEEE Int. Conf. on Robot and Human Interactive Communication (RO-MAN), 2024.
3. Luca Castri, Sariah Mghames, Marc Hanheide, and Nicola Bellotto, *Enhancing causal discovery from robot sensor data in dynamic scenarios*, Conference on Causal Learning and Reasoning, PMLR, 2023, pp. 243–258.
4. Luca Castri, Sariah Mghames, Marc Hanheide, and Nicola Bellotto, *Candoit: Causal discovery with observational and interventional data from time-series*, Advanced Intelligent Systems (2024).
5. Tomasz Piotr Kucner, Martin Magnusson, Sariah Mghames, Luigi Palmieri, Francesco Verdoja, et al., *Survey of maps of dynamics for mobile robots*, Int. Journal of Robotics Research (IJRR) **42** (2023), no. 11.
6. Timm Linder, Stefan Breuers, Bastian Leibe, and Kai O. Arras, *On Multi-Modal People Tracking from Mobile Platforms in Very Crowded and Dynamic Environments*, IEEE International Conference on Robotics and Automation (ICRA), 2016.
7. Timm Linder, Narunas Vaskevicius, Robert Schirmer, and Kai Oliver Arras, *Cross-modal analysis of human detection for robotics: An industrial case study*, Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS), 2021.
8. Sergi Molina, Anna Mannucci, Martin Magnusson, Daniel Adolfsson, Henrik Andreasson, et al., *The ILIAD safety stack: human-aware infrastructure-free navigation of industrial mobile robots*, IEEE Robotics & Automation Magazine (2023).
9. Olga Napolitano, *Shaping the information flow in robotic systems: from measures of information to active sensing control strategies*, Phd thesis, Uni. Pisa, 2024.
10. Olga Napolitano, Daniele Fontanelli, Lucia Pallottino, and Paolo Salaris, *Gramian-based optimal active sensing control under intermittent measurements*, 2021 IEEE International Conference on Robotics and Automation (ICRA) (2021).
11. Luigi Palmieri, Tomasz P Kucner, Martin Magnusson, Achim J Lilienthal, and Kai O Arras, *Kinodynamic motion planning on Gaussian mixture fields*, 2017 IEEE International Conference on Robotics and Automation (ICRA), IEEE, 2017.
12. István Sárándi, Timm Linder, Kai O. Arras, and Bastian Leibe, *MeTRAbs: metric-scale truncation-robust heatmaps for absolute 3D human pose estimation*, IEEE Transactions on Biometrics, Behavior, and Identity Science **3** (2021), no. 1, 16–30.
13. Bruno Siciliano, Oussama Khatib, and Torsten Kröger, *Springer handbook of robotics*, vol. 200, Springer, 2008.
14. Elisa Stefanini, Luigi Palmieri, Andrey Rudenko, Till Hielscher, Timm Linder, and Lucia Pallottino, *Efficient context-aware model predictive control for human-aware navigation*, IEEE Robotics and Automation Letters (2024).
15. David Taranta, Francisco Marques, et al., *An autonomous mobile robot navigation architecture for dynamic intralogistics*, IEEE 19th International Conference on Industrial Informatics (INDIN), IEEE, 2021, pp. 1–6.