Technische Universität München
TUM School of Computation, Information and Technology

Fleet Management for Mobile Service Robots Within the Operating Room Wing

Lukas Bernhard, M.Sc.

Vollständiger Abdruck der von der TUM School of Computation, Information and Technology der Technischen Universität München zur Erlangung eines

Doktors der Ingenieurwissenschaften (Dr.-Ing.)

genehmigten Dissertation.

Vorsitz: Prof. Dr. Viktor Leis

Prüfer der Dissertation:

1. Prof. Dr.-Ing. habil. Alois C. Knoll
2. Prof. Dr.-Ing. Johannes Fottner
3. apl. Prof. Dr. med. Dirk Wilhelm

Die Dissertation wurde am 02. April 2024 bei der Technischen Universität München eingereicht und durch die TUM School of Computation, Information and Technology am 17. Dezember 2024 angenommen.

Abstract

Motivation: Healthcare systems around the world are currently facing challenges due to increasing patient numbers, aging populations, and a shortage of qualified workers. To sustain and further improve patient care services, the application of assistive robotics is highly promising. For the first time, this doctoral thesis presents means for the dimensioning and composition of fleets of mobile service robots (MSR) within the operating room wing (OR wing) and introduces domain-specific concepts for orchestrating and managing them.

Contributions: For a context-dependent prioritization of competing task requests in critical or emergency-related situations, the novel scoring strategy *ANTS-OR* is proposed, which considers domain-specific information such as patient condition and clinical command hierarchy. The cost- and time-optimized assembly of individual task requests into overarching multi-task missions is governed by the *Vehicle Routing Problem for the Operating Room Wing* (VRP-OR), which is introduced and formally described for the first time. Due to NP-hardness, a heuristic algorithm is proposed for solving instances within acceptable processing time. This satisfies the requirements of the surgical application scenario, which may require the rapid adaptation of robot mission schedules to changing circumstances. To further improve robotic performance, a proactive behavior of MSR fleets is envisioned, by means of anticipating the use of sterilely packaged surgical instruments and materials in order to prepare likely-to-be-requested tasks in advance. For developing and evaluating the proposed concepts, the novel software framework *FleetOR* was implemented, which allows for the simulation of MSR fleets within OR wing environments. The simulation framework provides detailed models of the surgical environment and workflows based on extensive observations conducted in real-world OR wing facilities.

Results: Key findings include guidelines regarding the optimal dimensioning and composition of MSR fleets for the OR wing, taking into account robotic capabilities, driving speeds, battery capacity, workloads as well as limitations of the environment. An evaluation of the ANTS-OR task prioritization strategy showed clear benefits when dealing with critical situations, during which important tasks must be reliably identified and executed with priority. By means of multi-task mission planning, MSR fleet performance was shown to be highly improved, leading to a major reduction of driving durations and, effectively, reaction times. The proposed heuristic solver for the VRP-OR yielded close-to-optimal results while requiring extremely short computation times, proving its suitability for highly dynamic OR wing workflows. Further simulation results demonstrated that urgently needed surgical materials can be provided faster to the surgical team by means of anticipating future needs and holding likely needed materials readily available.

Conclusion: By means of novel fleet management techniques, important requirements of the surgical domain can be addressed and current limitations of MSR compared to human performance can be compensated without requiring improvements to robotic hardware. Focusing on aspects of integration and orchestration is a highly effective strategy for unlocking the unused potentials of mobile robotic technology and paving the way toward real-world application. When combining the insights gained in the course of this doctoral work, a performance level close to the human-only status quo can be achieved, which demonstrates that mobile service robotics is indeed a valuable technology for overcoming key challenges faced by today's surgical care services.

Zusammenfassung

Motivation: Gesundheitssysteme weltweit stehen derzeit vor Herausforderungen, die sich aus steigenden Patientenzahlen, einer Überalterung der Bevölkerung und einem Mangel an qualifiziertem Fachpersonal ergeben. Um die Patientenversorgung aufrechtzuerhalten und weiter zu verbessern, ist der Einsatz von Assistenzrobotik vielversprechend. In dieser Dissertation werden erstmals Konzepte zur Dimensionierung und Komposition von Flotten mobiler Serviceroboter (MSR) innerhalb des OP-Trakts vorgestellt und domänenspezifische Konzepte zu deren Orchestrierung eingeführt.

Konzepte: Die Bewertungsstrategie *ANTS-OR* ermöglicht eine kontextabhängige Priorisierung konkurrierender Tasks in notfallbezogenen Situationen. Dabei werden domänenspezifische Informationen wie der Zustand des Patienten oder die klinische Befehlshierarchie berücksichtigt. Die kosten- und zeitoptimierte Zusammenfassung einzelner Tasks zu übergreifenden Missionen wird durch das *Vehicle Routing Problem for the Operating Room Wing* (VRP-OR) beschrieben, das erstmals eingeführt und formal definiert wurde. Um Instanzen dieses NP-schweren Problems innerhalb akzeptabler Rechenzeiten lösen zu können, wurde ein heuristischer Algorithmus entwickelt. Dieser ermöglicht eine schnelle Reaktion auf sich änderende Umstände – eine zentrale Anforderung der chirurgischen Domäne. Durch ein proaktives Verhalten der Flotte kann die Leistungsfähigkeit weiter gesteigert werden, etwa durch eine Antizipation und Vorbereitung zukünftiger Aufgaben. Zur Evaluation der entwickelten Konzepte wurde das Software-Framework *FleetOR* implementiert, das die Simulation von MSR-Flotten basierend auf realitätsnahen Modellen der OP-Umgebung sowie aufgezeichneten chirurgischen Workflows ermöglicht.

Ergebnisse: Zu den zentralen Ergebnissen der Arbeit gehören Leitlinien für die optimale Dimensionierung und Zusammensetzung von MSR-Flotten für den OP-Trakt, wobei Roboterfähigkeiten, Fahrgeschwindigkeiten, Batteriekapazitäten, auftretende Arbeitslasten und besondere Anforderungen der Umgebung berücksichtigt wurden. Eine Evaluierung von ANTS-OR zeigte deutliche Vorteile bei der Bewältigung kritischer Situationen, in denen dringende Aufgaben zuverlässig identifiziert und mit Vorrang ausgeführt werden müssen. Durch die optimierte Zusammenfassung von Tasks zu Missionen konnten Fahr- und Reaktionszeiten der Roboter stark verbessert werden. Hierbei lieferte die entwickelte heuristische Lösungsstrategie für das VRP-OR annähernd optimale Ergebnisse bei gleichzeitig extrem kurzen Rechenzeiten, wodurch die Kompatibilität mit den hochdynamischen Prozessen des OP-Trakts gezeigt werden konnte. Weitere Simulationsergebnisse zeigten, dass benötigte chirurgische Materialien schneller zur Verfügung gestellt werden können, indem der zukünftige Bedarf antizipiert wird und wahrscheinlich benötigte Materialien situationsabhängig bereitgehalten werden.

Fazit: Mithilfe neuer Ansätze für Flotten-Dimensionierung und -Management können wichtige Anforderungen der chirurgischen Domäne erfüllt und die Limitierungen heutiger MSR im Vergleich zur menschlichen Leistung kompensiert werden, ohne dass hierfür Hardware-Verbesserungen nötig sind. Die Fokussierung auf Aspekte der Integration und Orchestrierung ist vielversprechend, um ungenutzte Potenziale der mobilen Robotik zu erschließen und den Weg in die reale Anwendung zu ebnen. Durch die Anwendung der gewonnenen Erkenntnisse kann ein Leistungsniveau erreicht werden, das nahe am rein menschlichen Status quo liegt. Dies zeigt, dass die mobile Servicerobotik in der Tat entscheidend zur Bewältigung der heutigen Herausforderungen in der chirurgischen Krankenversorgung beitragen kann.

Acknowledgement

I would like to express my sincere gratitude to my doctoral supervisor Prof. Dr.-Ing. habil. Alois C. Knoll (Head of the Chair of Robotics, Artificial Intelligence and Real-Time Systems at TUM) for his enduring support and valued guidance. Thank you for being a great inspiration and for sparking numerous new contacts, ideas and impulses. Likewise, I would like to express my sincere gratitude to Prof. Dr.-Ing. Johannes Fottner (Head of the Chair of Materials Handling, Material Flow, Logistics at TUM), who has co-supervised my doctorate and fostered an extensive exchange with research associates of his chair.

Furthermore, my heartfelt gratitude goes to Prof. Dr. med. Dirk Wilhelm, head of the research group MITI, who has more than exceeded his role as mentor to my doctorate. Thank you for carving valuable time out of your incredibly packed schedule, thank you for believing in my work and thank you for being a great friend. I would also like to express my sincere gratitude to Prof. Dr. med. Hubertus Feußner, the founder of MITI, whose infallible vision has sparked the ideas leading to the research project AURORA. Thank you for supporting me with your time and knowledge and for inspiring me with your forward-thinking attitude. Another big thank you goes to my amazing team at MITI. It was an immense pleasure working with you all and you have made these years truly special!

Moreover, I would like to thank the Bayerische Forschungsstiftung for funding the research project AURORA and thus making this doctorate possible. Thank you to all my colleagues of the AURORA team, with whom I have jointly worked towards the realization of the *Robotic Circulating Nurse* for a period of over three years. Together we have achieved great results, which have attracted international attention and have influenced this dissertation in many ways.

Of course, another heartfelt thank you goes to my parents, Sigrid Bernhard and Günter Möbus, who have always supported me both in good times and in bad. Thank you for your encouragement and for believing in me. Lastly, I would like to thank Johanna Rollenmiller, my long-time girlfriend and partner in crime, for *always* having my back. I simply could not have done it without you!

Munich, April 01, 2024 *Lukas Bernhard*

Contents

1

Introduction

1.1 Challenges in the Healthcare Sector

In the last two decades, the shortage of qualified staff has emerged as one of the most pressing challenges in healthcare [Mar+19]. This phenomenon can be observed in healthcare systems around the world, including Germany [Deu21], the United States of America [NSI23] and China [WZY16]. In 2021, 84 % of German hospitals reported that open nursing positions could not be fully staffed to meet their needs in human resources [Deu21]. Especially for larger institutions, this problem is omnipresent: Almost all (97 %) hospitals with more than 600 beds are experiencing severe difficulties in recruiting nursing staff [Deu21].

This is due to numerous factors, such as work environment, job satisfaction and low wages [Aik+01; Hus+12]. Studies indicate a high prevalence of distress in healthcare professionals, which often causes anxiety and depressive disorders [WC00]. Due to chronic understaffing, work conditions within hospitals are further worsened, since fewer persons are required to execute more tasks and care for more patients. This effectively creates a vicious cycle. Current wage levels cannot sufficiently compensate for such a constantly high mental and physical stress load, which causes workers to look for more attractive labor conditions elsewhere. Work-life balance and a healthy work experience become increasingly important concepts demanded by the young generation but are not yet broadly considered in healthcare management.

The advent of the Covid-19 pandemic has further worsened the situation, due to the extraordinarily high volume of patients that hospitals were facing throughout several peaks along the pandemic progress. A constant work overload combined with the risk of infection often caused physical exhaustion, fatigue, disorientation and even longer-term mental disorders like depression, anxiety, and post-traumatic stress disorder [Okp22; Bil+20]. As a result, this challenging conditions have caused further urgently needed workers to leave the healthcare sector.

Due to the acute staff shortage, clinical facilities can no longer be operated to the expected extent. This forces hospital managements to close staff-intensive units such as operating rooms (ORs) and intensive care units, while less personnel-dependent divisions, such as patient wards, are either downsized or operated at the absolute limit of available resources. As a result, an increasing number of surgical interventions must be canceled or postponed, since patients can no longer be appropriately accommodated and operating rooms can no longer be staffed. These conditions are not pessimistic visions of the future but are already becoming clearly apparent today. At the same time, the resulting limitations in healthcare provision have a direct impact on society and our quality of life. The standard of patient care, which has been excellent in countries of the developed world for many decades, is on the verge of becoming unsustainable.

Unfortunately, the prospects of improvement are poor: In the *Global strategy for human resources for health: workforce 2030* [Wor16], the World Health Organization projects an expected shortage of 18 million health workers by the year 2030. Longer life expectations and the ongoing inversion of the population pyramid observed for western societies lead to rising patient numbers and increasingly complex multi-disease cases. Growing case numbers are already reported throughout the medical fields, including cardiology [AV20], oncology [Bok+20] and diabetology [Deu22].

Solutions to these problems are long overdue. While societal and political reforms are clearly needed, the unused potentials of technological means such as digitalization, artificial intelligence and robotics must be exploited as well to fill in for lacking resources and improve the working conditions, while continuing to ensure a profitable operation of hospitals. Beyond addressing the immediate challenges, this is also a great opportunity to modernize structures and processes of the healthcare system, for not only maintaining the quality of care but continually increasing it. In the future, new demands will be placed on medicine: Treatment outcomes will need to be further improved in order to achieve higher cure rates, improved life expectancies and better quality of life. Patients will be increasingly placed at the center of the treatment process to improve patient experience and satisfaction [Ama+21]. Clinical services will need to be extended to provide care to more people and offer treatments for currently incurable diseases. Thus, wherever possible and ethically sound, processes must be automated, while preserving the crucial human-to-human relationship between patients and caregivers. Clinicians must be supported by assistive technologies to create the necessary conditions for them to fully focus on their strengths as human beings, while leaving dull, non-ergonomic and exhausting tasks to machines.

1.2 Unused Potentials

For being able to reach these ambitious goals and address the challenges described in the previous, available technologies must be leveraged and translated to the points of need. In this context, one promising and, as of yet, widely underutilized technology are mobile service robots (MSRs). According to the ISO 8373 standard released by the International Organization for Standardization, a service robot "*performs useful tasks for humans or equipment*", for which it requires a certain degree of autonomy, meaning the *“ability to perform intended tasks based on current state and sensing, without human intervention”* [Int21].

As shown in Figure 1.1, such systems typically consist of a mobile base, which is equipped depending on the intended duties of the robot. In the most basic case, only a load bed may be provided, while more complex systems may feature robotic manipulation arms or other end-effectors to interact with the environment. Since service robots usually need to communicate with human workers (e.g., to receive new task requests or to report their status) human-machine interfaces, such as displays or voice interaction systems, are also important building blocks.

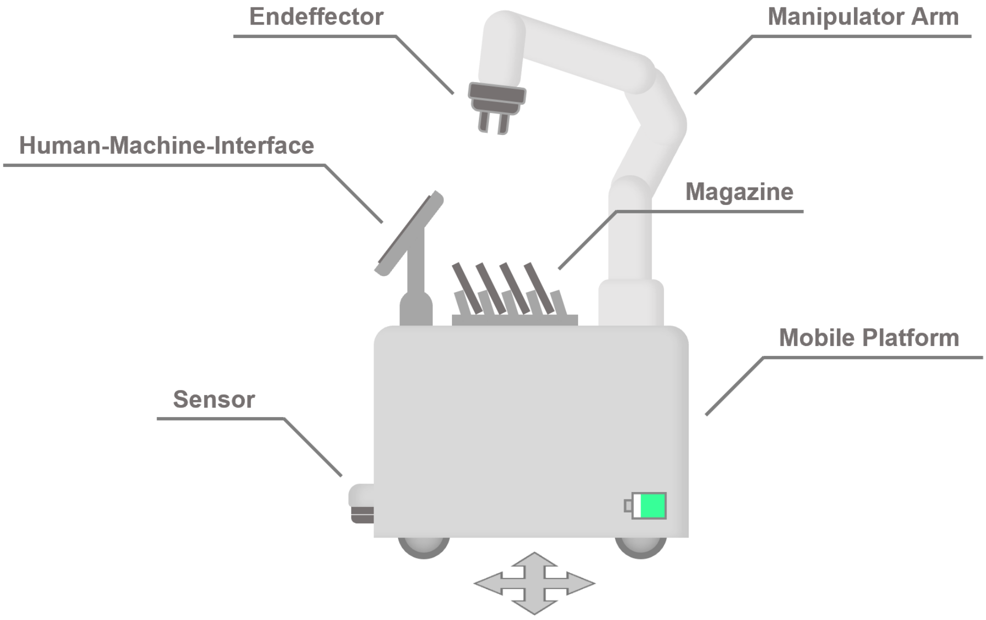

Figure 1.1: A conceptual illustration of an MSR is shown. By means of a mobile platform, the robot is able to move within the environment. For perceiving obstacles and objects of interest, the robot is equipped with sensors, such as depth cameras or *light detection and ranging* sensors (LIDAR). The robot may be able to physically interact with its environment using one or more robotic manipulator arms and task-specific endeffectors. For storing objects, the robot may possess a magazine or inventory mounted on its mobile platform. Via human-machine-interfaces, such as touch displays or voice control systems, the robot receives requests from the user and reports its status.

MSR technology has advanced considerably over the recent decades and has found real-world applications in several domains, such as delivery, manufacturing and intra-logistics. In the hospital context, a variety of use cases have been proposed in academic literature and evaluated on a proof-of-concept basis. Most work focuses on the patient ward and intra-logistical tasks in non-public areas of the hospital. However, the transition from research prototypes to real-world products has been very limited so far. Other clinical environments have been neglected almost entirely: Personnel-intensive units, such as operating room wings (OR wings) in particular, are heavily affected by the staff shortage problem and can only function with sufficient resources. At the same time, there is great potential for mobile robotic automation: Tasks such as the transportation of objects or the adjustment of devices must be carried out frequently when preparing and conducting surgical procedures and are well-suited for robotic execution. This is explored by the research project *Auto-navigating Robotic Operating Room Assistance* (AURORA) [Ber+22], which currently spearheads the development of MSR technology for the OR wing (see Figure 1.2) and also provides the framework for this doctoral thesis. AURORA focuses on two use cases within non-sterile parts of the OR wing:

1. **Provision of sterile goods**: Here, the AURORA robot autonomously navigates to the storage room, retrieves the requested articles from the magazine, moves to the operating room and approaches a hand-over location near the patient table, while keeping a safe distance to all sterile surfaces. The packaging of the material is now opened by the robot and handed over, such that the content can be hygienically retrieved by a sterilely dressed member of the surgical team.
2. **Adjustment of medical devices**: For that, the AURORA robot moves to the medical device and adjusts it according to the instructions of the surgical team. For that, it interacts with the haptic interfaces typically provided on the front of the device, which would normally be operated by humans. Hereby, the use of digital input interfaces is avoided, which may be difficult to realize in the near future due to regulatory hurdles.

While the capabilities of the AURORA robot will not yet fully match those of human workers, it will relieve the OR staff from frequent and monotonous tasks. Thereby, the overall workload is reduced and OR nurses are enabled to focus on other tasks, which may be more fulfilling and in line with their strengths as human beings. This will contribute to making OR nursing jobs more attractive again. With the help of MSRs such as AURORA, surgical procedures can be conducted with less personnel, since tasks can be split between robots and humans. This allows for keeping more operating rooms in service and, at the end of the day, treating more patients.

However, the integration of MSRs into the OR wing is challenging due to the high safety requirements and the complex and dynamic nature of the environment. In such settings, robotic systems are required to drive slowly and navigate carefully around dynamic objects, moving persons and sterile surfaces. As of yet, it is unclear how severely these circumstances limit the effectiveness of singular robotic systems operating independently and whether further improving robotic hardware can alleviate this. In order to exploit the full potential of MSR technology for the OR wing and reach an acceptable performance level, the formation of robotic teams that are globally governed by a fleet management system is promising. This approach is expected to increase the total robotic workforce, allow for a better parallelization of tasks and compensate for shortcomings of individual systems. By means of global orches-

Figure 1.2: A conceptual rendering of the AURORA robot (left side) within the OR environment is shown. AURORA supports the surgical team by providing surgical materials and adjusting devices in the immediate surroundings.

tration, it should become possible to flexibly reassign robotic resources to locations where they are currently needed most in order to collaboratively handle situations of high demand.

1.3 Problem Statement

This doctoral thesis is concerned with describing the unique characteristics and challenges of the OR wing environment with regard to the integration of MSR fleets. This includes fundamental investigations regarding the dimensioning, composition and management of MSR fleets, which will then be used for developing approaches for exploiting the full potential of this technology and overcoming its current limitations. A central aim of this thesis is to demonstrate the feasibility and effectiveness of MSR technology for OR wing applications and provide guiding directions for moving forward with real-world translation.

As of yet, the OR wing has not been studied with regard to the integration of individual MSRs, let alone MSR fleets. In order to make full use of the expected benefits, the target environment must be thoroughly understood and, if necessary, carefully adapted to the needs of mobile robotics. It is still entirely unclear, if and how MSR technology can be applied effectively to support workflows within the OR wing. This includes fundamental aspects such as the number of robots required to handle everyday workloads and whether the resulting fleet size is at all reconcilable with space limitations and economic constraints. It is further to be determined which robot velocities are necessary to achieve the necessary performance level and whether these can be safely achieved within the target environment. The composition of fleets from robots with different capabilities is another fundamental aspect to be studied and will yield implications for the future design of robotic systems.

Beyond these elementary insights, new concepts for fleet management are lacking that fit the requirements of the OR wing, e.g. regarding the prioritization of tasks in situations with high demand and conflicting interests, the planning of multi-task missions along points-of-interest of the target environment and the improvement of the material flow from storages to the final users. In order to compensate for shortcomings in robotic performance that may occur, the fleet management should be able to anticipate upcoming user needs, in order to prepare in advance and thus ensure a smooth and efficient surgical workflow without causes of frustration for the surgical teams.

In summary, the following central research questions were identified for this doctoral thesis, aiming at laying the groundwork for integrating MSR fleets into OR wing environments:

Which unique characteristics and requirements must be addressed when integrating MSRs into the OR wing?

How must MSR fleets be dimensioned, composed and managed to optimize the performance for the OR wing application scenario?

How should processes and environmental features be modified to support this?

What performance levels can be achieved compared to the human-only status quo? Is this sufficient for real-world needs?

1.4 Summary of Contributions

The following section briefly outlines the methods used to investigate the research questions stated above, followed by a summary and interpretation of the obtained results. Finally, central aspects of the presented work are discussed with respect to the interests of key stakeholders.

Methods

As a means for studying different fleet configurations and validating the developed concepts, a simulation-based approach was used. For that, the novel software framework FleetOR (section 4.3) was developed, which provides the means for simulating MSR fleets within OR wing environments, including exposure to realistic workflows and workloads. A model OR wing was created based on real-world facilities at a German university hospital (section 4.2.2). By means of observational studies, the workflow of surgical interventions was recorded with respect to potential tasks suitable for mobile robotic execution (section 4.2.1). As a further core part of the simulation framework, means for freely assembling MSR fleets were implemented. This allowed for studying fleets of different sizes and compositions (section 4.3.5).

In order to address the domain-specific needs of the OR wing, novel fleet management concepts were developed. To handle critical or emergency-related situations with varying task priorities, the novel scoring strategy ANTS-OR was introduced (section 4.4.2). This prioritization approach specifically considers domain knowledge such as patient condition and clinical command hierarchy. The cost- and time-optimized assembly of individual task requests into overarching multi-task missions is governed by the *Vehicle Routing Problem for the Operating Room Wing* (VRP-OR), which, as a key contribution of this doctoral thesis, was motivated and formally described for the first time (section 4.4.3). Due to the NP-hardness of this problem, a heuristic algorithm was proposed to solve instances within acceptable processing time. This allows for a rapid adaptation of robot mission schedules to changing circumstances, which is a central requirement of the surgical domain. To further improve robotic performance, a proactive behavior of MSR fleets was envisioned, where the intraoperative use of sterilely packaged surgical materials is anticipated to prepare likely requests in advance.

Results

As a fundamental question in the context of MSR fleets for the OR wing, the adequate dimensioning of fleets was investigated (section 5.1). This was used for providing recommendations regarding the optimal number of robots per operating room, necessary for handling realistic workloads. The optimal fleet size was found to heavily depend on robot driving speeds: while 1 robot per OR suffices for fast driving speeds ($1.2m/s$), 3-6 robots are required for slower driving speeds ($0.3m/s$). Since high driving speeds are difficult to realize within the complex, dynamic and safety-critical OR wing environment, means for circumventing this problem were studied (section 5.3). By introducing zones of different driving speeds, an acceptable compromise between driving speed and fleet performance was achieved. A further key insight was the distinct superiority of allrounder robots over specialized robots with respect to fleet performance (section 5.2).

The merits of the ANTS-OR task prioritization strategy were evaluated and compared to a first-in-first-out approach (section 5.5). The results show a clear benefit when dealing with emergency situations: Using ANTS-OR, critical tasks are executed considerably sooner than less important tasks associated with a lower priority. The same was shown for tasks originating from staff members with a high rank within the clinical command hierarchy.

When evaluating a novel greedy mission planning algorithm proposed for heuristically solving instances of the VRP-OR, a considerable reduction of task driving times was shown for multi-task missions, as compared to executing tasks in an isolated fashion (section 5.6). Even for a relatively short mission length of 4 tasks, driving durations were reduced by 48.7 % on average. Only a rather small average increase in driving duration (6.8 %) was observed when comparing the results of the heuristic solver to the optimal solution. At the same time, the observed computation times were extremely short (<38.9 ms).

In order to further improve fleet performance, the logistical concepts of quick storages (section 5.7) and robot inventories (section 5.8) were introduced. Both approaches anticipate materials that are likely required during a given surgery. Especially the introduction of quick storages has shown great potential for reducing task execution durations, even for lower material prediction accuracies.

Lastly, all previous findings and concepts were combined into a final approach and benchmarked against both a human-only scenario and a "naive" robotic reference scenario (section 5.9). Only by incorporating the results of this doctoral thesis, an acceptable performance of the robotic fleet was reached, when comparing to the human benchmark. The outcome of the naive reference scenario was inferior by one order of magnitude and proved to not be sufficient for real-world requirements.

Conclusions

The obtained results strongly indicate that the integration of MSR fleets into surgical OR wing environments is indeed feasible and that real-world workloads can acceptably be handled when employing adequate techniques for fleet dimensioning and fleet management. An orchestrated, time-optimized and proactive fleet behavior was shown to be highly beneficial for alleviating the drawbacks of individual robotic fleet members regarding driving and manipulation speeds. At the same time, essential domain-specific requirements were addressed, ensuring compatibility with real-world surgical workflows that may involve highly dynamic situations and critical patient conditions. The idea of globally managing and optimizing OR wing processes on a fine-granular level introduces a fundamental paradigm change compared to the established approach of managing operating rooms individually while coordinating processes only on a high level.

Table 1.1: The concepts presented in this doctoral thesis benefit multiple stakeholders within the healthcare sector (also see chapter 6). These relationships are summarized in the following table.

	Clinical Staff	**Patient**	**Clinic / Hospital**	**Healthcare System**
Fleet Dimensioning	✔		✔	✔
Speed Management	✔	✔		
Battery Dimensioning	✔	✔	✔	✔
Task Prioritization	✔	✔		
Multi-Task Missions	✔	✔	✔	✔
Quick Storages	✔	✔	✔	✔
Robot Inventories	✔	✔	✔	✔

When taking the perspectives of key stakeholders such as clinical personnel, patient, hospital and healthcare system, it can be concluded that all parties benefit from the application of globally orchestrated MSR technology (see Table 1.1). Clinicians benefit from a compensation of lacking personnel and by a relief from dull or non-ergonomic tasks, which helps to create a better work environment and thus prevent stress and burn-out. Patients benefit from caregivers having more time for providing direct care and from an improved capacity of surgical care services, resulting in fewer postponements and cancellations of planned surgeries. Hospitals benefit from being able to offer more attractive jobs, improve the long-term retention of personnel and increase revenues by treating more patients. The healthcare system overall benefits from an increased robustness of patient care services even in times of high demand.

2

State of the Art

Various applications of mobile robotics in the hospital have been proposed in scholarly literature and a limited number of systems are already in operation within real-world facilities. The range of proposed use cases, along with the presented prototypes or products are summarized in section 2.1. For some of these application scenarios, the subject of mobile robotic fleet management has been studied as well, which is covered by section 2.2.

2.1 Mobile Robotics for the Hospital

In the hospital context, possible applications of MSRs are manifold and numerous approaches have been proposed in recent decades. The technology is envisioned to be an important puzzle piece of the hospital of the future [Ama+21; Ama+22]. However, the challenges are enormous, due to the heterogeneity and complexity of clinical environments [Guz+21]. Nowadays, successful translation to the clinic has been accomplished for logistical robotic fleets within some clinical environments. Other application fields, such as rehabilitation and care, have been explored mostly in the academic context. An overview of research projects and commercially available systems is given in the following.

2.1.1 Intralogistics

As of yet, the majority of MSR concepts presented in industry and research can be categorized as intralogistical robots used for executing transportation tasks within hospital environments. These systems often take the form of *Automated Guided Vehicles* (AGV), which are motorized transportation platforms equipped with sensors for perceiving their environment. AGVs are usually monitored and organized in fleets by connecting them to a central management software. Nowadays, the use of AGVs has become quite well-established in industrial manufacturing environments and is also gaining ground in the healthcare domain. More advanced systems may also be equipped with means to interact with the environment, i.e. by using a robotic manipulator mounted on the mobile platform to extract materials from a storage cabinet. According to a case study by Ozkil et al., hospitals can greatly benefit from the implementation of such systems, which allows for the automation of workflows that are currently

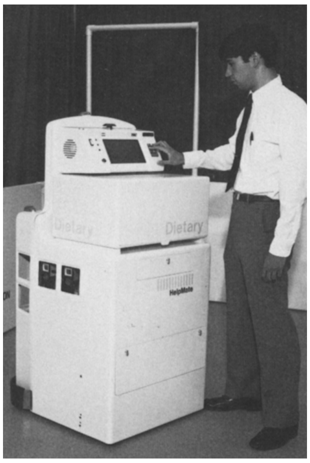

(a) HelpMate (Evans et al.) [Eva+89]

(b) T3 (Aethon) [Aet24]

(c) Moxi (Diligent Robotics) [Dil24]

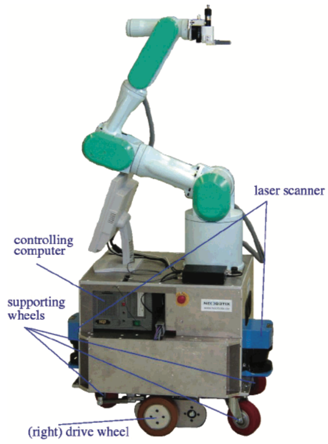

(d) Laboratory Robot (Knoll et al.) [Kno+04]

(e) ROGER (Gross et al.) [Mül+20][1]

(f) Penelope (Kochan) [Koc05]

Figure 2.1: Exemplary robotic systems from related work are shown.

heavily dependent on manual transportation [Ozk+09].

Most MSR systems for hospital logistics are designed for the patient ward or for off-stage facilities related to material supply, warehousing and housekeeping. In scientific literature, various approaches have been presented, ranging from general-purpose transportation applications [Eva+89; KE92; Eva94; Fun+03; Tak+09; Wan+09; Tak+10; Lju+12; Tak+12; Kha13; SSS13; AMN15; Bac+17; PGP17] (see Figure 2.1a) to the transportation of specific types of goods, such as food [Car+06; PGP17], medicine [BMC09; KRL09], laundry [PGP17], waste [PGP17], lab samples [Mur+09] and hospital beds [Wan+15; Wan+16; SW17].

Multiple commercial systems are available on the market today and are in use within healthcare facilities around the world. This includes general purpose transportation robots such as the *T3* (ST Engineering Aethon, Pittsburgh, USA), that has been deployed to over 100 hospitals [Blo11; Aet24] (see Figure 2.1b); the *MEDI MOVE* (ek robotics, Hamburg, Germany) [ek 24]; the *QC Bot Open* and *QC Lift* (Vecna Technologies, Burlington, USA) [Med24a; Med24b]; the *RobCab* (Robotdalen, Västerås, Sweden) [New14]; the *TIAGo Delivery* (PAL Robotics, Barcelona, Spain) [PAL22]; and solutions by Robotic Automation (Sydney, Australia) that have been deployed to Australian hospitals [Rob23]. Commercial systems specializing in the transportation of medications and lab samples include the *Relay* robot (Swisslog, Buchs, Switzerland) [Swi24] and the *HOSPI* robot (Panasonic, Kadoma, Japan) [Pan15].

2.1.2 Patient Care

Clinical and non-clinical patient care is another major application scenario for MSRs [NSN09]. A variety of concepts for the hospital have been presented in the scientific literature, including robots for monitoring patient conditions during patrol rounds [SIM03; BLX08], measurement of vital signs or other health-related parameters with and without physical contact [Gar+12; Ahn+14; Bro+15; Zuk+18; Cha+21; Wan+21; Hua+22], guidance of patients [SHC04; BLX08; Vas+17; MMD21], delivery of supplies to patients (such as sheets, clothes etc.) [SHC04], cleaning of patient rooms [BLX08], patient lifting and transfer [Muk+10; Hu+11; Ngu+13; Din+14], provision of information to patients [Ahn+14], assistance of patients in daily tasks [Cre+16], entertainment and social interaction with patients [Nej+07; AN08; Ahn+14], assistance of clinicians during ward rounds [Ili+14; Ter+14; Tas+15; Gar+16; Cro+18] and general-purpose concepts [Yuk+04; Nam+22]. A commercially available product is the *Moxi* robot (Diligent Robotics, Austin, USA) [Dil24] that aims at supporting the staff in clinical wards with non-patient-facing tasks, mainly the transport of materials (see Figure 2.1c). However, due to the avoidance of direct patient interaction, this system arguably leans more towards being a logistical robot – although acting within a care environment.

2.1.3 Surgery

While there is great potential for easing staff workload and overcoming personnel shortages, the technological and environmental barriers for integrating MSR technology into the OR or the OR wing are considerable. Due to that, research and development efforts have been very limited so far, with the exception of the aforementioned project *AURORA* [Ber+22; Bau+22;

[1]licensed under https://creativecommons.org/licenses/by/4.0/

MBW22; Ber+24], which provides the framework for this doctoral thesis. AURORA aims to design and implement a self-navigating robotic assistance system for executing tasks within the non-sterile part of surgical operating rooms. Central use cases include the fetching and handing over of sterile materials as well as the operating of medical devices via available human-machine interfaces. One or even multiple robots may be teamed up with human circulating nurses to collaboratively assist in concurrent surgical interventions within the OR wing.

Figure 2.2: The AURORA prototype features a mobile omnidirectional base platform equipped with a robotic manipulator arm. A custom-designed endeffector and an auxiliary mechatronic module enable the robot to open sterilely packaged surgical materials.

The AURORA prototype is shown in Figure 2.2. It features a custom-designed mobile base platform that enables the robot to move omnidirectionally. Since space inside operating rooms is usually at a premium and sterile boundaries must be respected, the capability to move sideways and turn on the spot is highly beneficial. Using its mobile platform, the robot is able to move within non-sterile spaces and reach points of interest inside and outside of operating rooms, such as storage cabinets, medical devices, disposals etc.

The platform is equipped with the 7-DOF robotic manipulator arm *Panda* (Franka Emika, Munich, Germany), which features a custom-built endeffector based on the Panda's standard gripper. The packaging of sterile surgical material typically features two flaps by which it can be pulled open in order to reveal the content. Using its endeffector, AURORA can grasp these flaps and pull them apart. This is achieved using an auxiliary mechatronic module mounted on the mobile platform, which fixates the material and spreads the flaps such that the endeffector can grip them. After opening the packaging, the material is handed over to a member of the surgical team, typically the scrub nurse. This is achieved using a gesture-based interaction routine (see Figure 2.3), which is initiated by an upward pointing index finger, followed by pointing to a suitable handover position. The robot now delivers the opened material to the hand of the user, such that the sterile content can be retrieved in a hygienic manner, i.e. without touching parts of the robot or the non-sterile outer surface of the packaging. For detecting the gestures, a *Realsense D435i* camera (Intel Corporation, Santa Clara, United States) is used, which is mounted on top of the endeffector.

Using the tip of its endeffector, AURORA is able to physically interact with haptic interfaces of medical devices, such as buttons and touch screens. For that, ArUco markers are attached to the devices, such that the robot can detect the position of control elements using the camera. The current state of a device can either be read using a digital interface or by directly observing it from the display using image recognition. By interacting with the control elements, the robot can modify the device state according to the needs of the user. This is further described in [Bau+22].

While AURORA is the first *mobile* service robot for the OR wing, other assistive robots have been presented in scientific literature that are related to the OR circulator use case. A variety of existing work envisions a *Robotic Scrub Nurse* (RSN), aiming at the automated handover of sterile surgical instruments. Such systems are typically statically mounted within the sterile zone of the operating room and thus must comply with sterility requirements. The core task of these systems is the management of a selection of surgical instruments that must be handed over to the surgeon on request or following a predictive strategy. After the introduction of the Penelope system by [Koc05] (see Figure 2.1f), several further approaches have been proposed, including [Tak+08; Car+09; Yos+10; Jac+12; Per+12; Kog+19; She+19; Wag+23b]. Similar to the roles of human scrub and circulating nurses, the concept of the RSN can be seen as complementary to AURORA, with one robot acting within the sterile zone and the other one acting within the non-sterile surroundings. For future development efforts, an interaction between both sides can be envisioned, leading to a fully automated workflow for the provision of sterilely packaged surgical materials: A desired material is collected, transported, and opened by AURORA and subsequently handed over to an RSN system, which further prepares the material, before handing it over to the surgeon.

2.1.4 Laboratory Automation

In [Lüt+00; Sch+03; Kno+04; Wes+04; Woj+07], Knoll et al. successfully demonstrate the use of a mobile robot for sample management in biotechnical laboratories (see Figure 2.1d). The robot – which is based on the mobile platform *MP-L655* (Neobotix, Heilbronn, Germany) and is equipped with a *PA-10* robotic arm (Mitsubishi Heavy Industries, Tokyo, Japan) – is able to independently perform the required steps for monotoring cell cultures and determining the optimal harvest time. This involves the handling of samples and vials, the safe navigation of the laboratory environment and the robust operation of laboratory devices. While these are challenging tasks for robotic automation, usually requiring extensive human involvement, the authors are able to repeatedly demonstrate the successful use of the system for automated cell cultivation procedures. Moreover, the solution does not require any modifications of the standard laboratory equipment. Many concepts presented by Knoll et al. later proved to be highly useful for other application scenarios, for example in the context of the aforementioned research project AURORA, where the robust interaction between robot and medical devices is one central challenge.

(a)

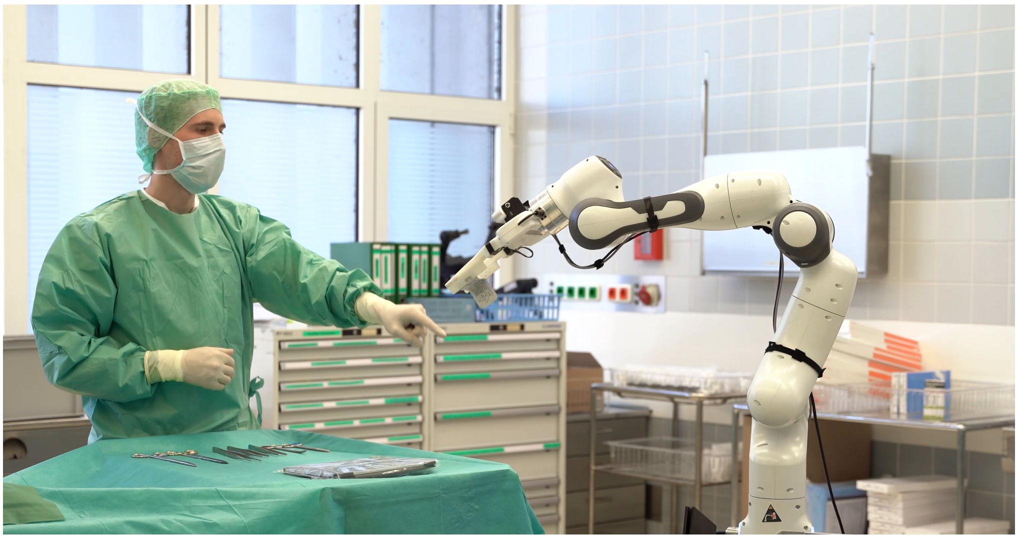

(b)

(c)

Figure 2.3: Using a gesture-based routine, clinicians can interact with AURORA in order to hygienically receive surgical materials that were procured and opened by the robot. (a) A raised index finger initiates the routine. (b) After pointing to the desired handover position, the robot delivers the material to the hand of the user. (c) The sterile content can now be carefully retrieved in a hygienic manner.

2.1.5 Rehabilitation

Gross et al. [Gro+14; Eis+15; Wen+16; Gro+17a; Gro+17b; Vor+18; Sch+19; Mül+20] propose to utilize MSRs for supporting inpatients during rehabilitation. The authors present their autonomous socially-assistive robots *ROREAS* and *ROGER* for assisting stroke patients during walking exercises (see Figure 2.1e). These robots are able to follow patients, give them motivational support and guide them back to their room, if need be. The authors demonstrate that their systems are able to navigate in crowded corridors within a rehabilitation clinic and reliably re-identify the patient even after temporal occlusion.

A study conducted by Röhner et al. (in cooperation with Gross and his working group), which used a similar system based on the robot platform *Tory* (MetraLabs, Ilmenau, Germany), indicates that robot-based gait training decreases a patient's length of stay in acute clinics [Röh+21].

2.1.6 Hospital Hygiene

In [CM14], Cepolina et al. present an MSR for cleaning and sanitizing hospital environments. By means of vacuum cups and linear actuators, the robot is able to climb walls and successively disinfect the traversed surfaces.

In the wake of the SARS-CoV-2 pandemic, further solutions for disinfecting hospital environments have been developed. Most systems consist of ultraviolet (UV) lamps mounted on a mobile driving platform, thereby being able to move through the hospital environment while applying UV-C type light (100–280 nm wavelength) to potentially contaminated surfaces. Multiple systems have been presented in academia [GGH21; NMJ21; Lor+22] and industry [Aka24; UVC23; UVD24; Xen23].

Another solution by *XDBOT* [Dem20] uses a nozzle to spray disinfectant onto potentially contaminated surfaces. By means of electro-charging the droplets, surfaces that have already been covered are repelling excessive disinfectant, pushing it towards untreated surfaces.

2.1.7 Social Assistance and Education

MSR technology also shows potential for implementing social assistants for the hospital. In scholarly literature, socially assistive robots have been presented that interact with inpatients or visitors, e.g. *HealthBot* by Jayawardena et al. [Jay+16] or *Brian* by Nejat et al. [NF08; AN08; Nej+07]. The humanoid MSR *Arash* by Meghdari et al. [Meg+18] supports children with cancer by providing assistance, education and entertainment.

Other work primarily focuses on offering educational services by means of telepresence such as the robots presented by Soares et al. [SKC17] and Schmucker et al. [Sch+20]. These systems aim at maintaining daily school routines and social interaction with classmates.

2.2 Robotic Fleet Management

MSR fleet management comprises several tasks with differing underlying problems. Firstly, the fleet must be dimensioned according to the imposed workload, while finding a suitable trade-off with economic aspects. In cases where different tasks require robots with differing capabilities, an optimal fleet composition must be found. This may be further complicated by robot types with overlapping capabilities. Moreover, it may be required to dynamically re-balance size and composition over time, in cases where demand changes drastically, or where fleet members are malfunctioning. As of yet, these problems are widely understudied for hospital-related use cases, with very few exceptions, as summarized throughout the following subsections.

A further key aspect of MSR fleet management is mission planning, which involves the allocation of one or multiple tasks to individual fleet members, while optimizing global cost values. Multi-robot task allocation (MRTA) has been studied extensively in the scientific literature and has sparked the progressive development of taxonomies for classifying MRTA problems appearing in different application scenarios. One renowned taxonomy was proposed by Gerkey and Mataric [GM04], who classify MRTA problems along the following axes:

- **Single-task robots (ST)** vs. **multi-task robots (MT)**: A given robot can either execute one task at a time (ST) or multiple tasks (MT).
- **Single-robot tasks (SR)** vs. **multi-robot tasks (MR)**: The execution of a given task may either require just a single robot (SR) or multiple robots (MR).
- **Instantaneous assignment (IA)** vs. **time-extended assignment (TA)**: Tasks are either assigned while considering future tasks (TA) or without considering future tasks (IA).

Using this taxonomy, problem classes can be denoted by concatenating the abbreviations, e.g., the most basic case of *single-task robots, single-robot tasks and instantaneous assignment* would be denoted as *ST-SR-IA*. Extending the work of Gerkey and Mataric, Korsah et al. [KSD13] propose their taxonomy *iTax*, which is now also applicable to problems with interrelated utilities and constraints. For that, the following four additional categories are introduced:

- **No dependencies (ND)**: The effective utility of a robot for a task is independent of other tasks or robots.
- **In-Schedule dependencies (ID)**: The effective utility of a robot for a task depends on other tasks in its schedule.
- **Cross-Schedule dependencies (XD)**: The effective utility of a robot for a task depends on its own schedule and the schedule of other robots.
- **Complex dependencies (CD)**: Like XD, however, tasks may be multiply decomposable (can be divided into sub-tasks in different ways) and the effective utility depends on the chosen decomposition.

Most notably, Korsah et al. state that multi-robot routing problems, which can be modeled as *multiple traveling salesman problems* (m-TSP) and are common in robotics, are now included in the new taxonomy.

A further taxonomy is proposed by Nunes et al. [Nun+17], which is also based on the work of Gerkey and Mataric and puts an emphasis on problems with temporal and ordering constraints. For that, a further axis is introduced: **Time Window (TW) vs. Synchronization and Precedence (SP) constraints**. The authors further distinguish between hard temporal constraints (no violations) vs. soft temporal constraints (violation with penalty) and deterministic models (output is determined by initial conditions) vs. stochastic models (subject to uncertainty).

Many MRTA problems are related to the well-known combinatorial *Vehicle Routing Problem* (VRP), where least-cost routes must be found for delivering goods from a central depot to customers using multiple vehicles [WN14]. Many variants of this problem have been described and studied in scholarly literature, such as the *Capacitated Vehicle Routing Problem* (CVRP) [SW19], the *Multi-Depot Vehicle Routing Problem* (MDVRP) [RGP20], the *Open Vehicle Routing Problem* (OVRP) [SW19] or the *Vehicle Routing Problem with Pickups and Deliveries* (VRPPD) [Ber+07]. The mission planning problem presented in this doctoral thesis can also be understood as a variant of the VRP, as elaborated in section 4.4.3.

Furthermore, fleet management solutions are either centralized or distributed [BXD10]. In a centralized architecture, all relevant information – such as robot states, robot capabilities, pending task requests, state of the environment etc. – is collectively made available to a singular fleet manager. An advantage of this approach is that a global perspective is taken, which facilitates the finding of globally optimal or close-to-optimal solutions. On the downside, large data quantities must be processed, which may lead to long computation times. Moreover, the entire fleet is effectively rendered inoperative in case the central fleet manager becomes unavailable, e.g., due to a technical problem. In a distributed architecture, the fleet is not managed by a singular entity. Instead, the problem is shared between multiple entities that do not necessarily possess all available information. Advantages of the decentralized approach are a reduced computational effort due to smaller problem instances and a higher robustness since other entities can potentially take over the duties of an inoperative agent. On the other hand, an optimal solution may only be found in a local, not a global sense.

2.2.1 Clinical Applications

In the following, related work concerned with MSR fleet management in the context of different hospital environments is presented. A delimitation of these approaches with regard to each other and the scope of this doctoral thesis is given in section 3.4.

Rossetti et al.

A first contribution to the field has been made by Rossetti et al. [RKF98], who propose a modeling approach for the performance evaluation of a homogeneous fleet of AGV-style laboratory delivery robots for the hospital. By means of simulation, the authors are able to show the superiority of an aptly-sized robotic fleet regarding costs, turn-around time, delivery variability and cycle time when compared to a team of human couriers – which then was (and mostly still is) the clinical state-of-practice.

Jeon et al.

Jeon et al. [JL16a; JL16c; JL16b; JLK17; JL17a; JL17b] have conducted work concerning the centralized real-time allocation of tasks to a homogeneous robotic fleet executing logistical tasks in a clinical environment. With reference to the taxonomy introduced by Gerkey and Mataric [GM04], the authors distinguish between *single-task allocation*, where only one task may be assigned to a robot, and *multi-task allocation*, where multiple tasks may be assigned to a single robot, e.g. for delivering multiple goods to different locations. For solving the multi-task allocation problem, which clearly is the more challenging of the two, the authors propose two approaches: A first-order search algorithm, which is limited to two tasks that can be assigned simultaneously, and a combinatorial search algorithm, which yields an improved schedule and has no limitations regarding the numbers of task assignments to a single robot. The authors also aim at finding the optimal fleet size for a given scenario, since – after a certain cut-off point has been reached – adding more robots to the fleet will increase costs without further improving overall throughput [JLK17].

Kumar et al.

Kumar et al. [KSW18] have proposed concepts for the centralized task allocation of homogeneous fleets consisting of logistical robots operating within the hospital. The authors aim to minimize the total movement distance of all robots. Two different scheduling schemes have been presented and evaluated: the deep Hungarian scheme, which is designed for scenarios where the location of jobs is known beforehand, and the deep Voronoi scheme, which does not require prior knowledge of jobs. The authors have been able to show that the traveled distance can be efficiently reduced using the algorithms. However, due to the NP-hardness of the problem, only a sub-optimal solution can be found within reasonable time.

IWARD

Multiple approaches for the scheduling of (semi-)heterogeneous fleets of multi-purpose robots within the clinical ward have been proposed in the context of the research project IWARD (Balbaaki et al. [BLX08; BX09; BXD10; BX10], Thiel et al. [THB09] and Shi et al. [Shi+10]). The robots developed by the authors are based on identical mobile platforms but can be equipped with additional modules in a plug-and-play fashion. Depending on the individual configuration, the robots are intended to execute different missions, including delivery of goods, cleaning, guidance of patients/visitors, patient monitoring, patient-doctor teleconferencing, hospital patrolling and environmental monitoring. The workgroup reports the prototypical development of modules for patient/visitor guidance [HHS10] and patient monitoring [Mam+14].

In [BLX08], Baalbaki et al. introduce the problem of simultaneously determining both the optimal configuration of a given robot within the fleet (meaning the type of module installed on the base platform) and its location (meaning the assigned home stations of the robot, which function as start and end points of task executions). This information is used for planning upcoming shifts, which can thus be classified as an offline scheduling problem without real-time requirements. To solve this problem, the authors propose a linear programming model and the usage of a column generation approach for obtaining near-optimal solutions. In [BX09], this concept is extended to include real-time mission assignment and power management. The authors conclude that their implementations are only feasible for small instances of the problem due to computational complexity. In [BXD10] Baalbaki et al. extend their work to include real-time mission assignment and scheduling based on evolutionary algorithms [Sch94; Pri04]. An updated version of the schedule is calculated every time new mission requests arrive or the composition of the fleet changes. Several evolutionary algorithms are proposed and benchmarked that provide sub-optimal solutions in comparatively short processing durations (approx. 30s for 7 robots, 49 missions and 4 functionality modules).

In [BX10], Baalbaki et al. propose a distributed scheduling approach, as an alternative to the centralized concepts presented so far: Based on the dynamic task allocation system MURDOCH [GM00; GM02] and the Contract Network Protocol [DS83], the role of the controller is introduced, which is assigned to exactly one member of the fleet at a given point in time. The controller broadcasts incoming missions to all members of the fleet, which use an internal module named decision finder to calculate an estimated schedule and return a calculated penalty for the execution of the task. Finally, the task is assigned to the fleet member with the lowest penalty value. Baalbaki et al. numerically benchmark different configurations of their distributed scheduling mechanism. The authors conclude that their rescheduling techniques provide a significant benefit regarding the redistribution of workload and the reduction of delays within the system. Thiel et al. contribute further general considerations regarding robust mission handling and multi-robot missions [THB09], however, these aspects were not included in the implementations of Baalbaki et al. [BX10].

Guzman et al.

The *ENDORSE* research project aims at facilitating the deployment of MSRs into health centers [Guz+21]. In this context, Guzman et al. [Guz+21] present a centralized fleet management system for MSR fleets within real-world hospital environments. The authors focus on managing the execution of logistical tasks carried out within the ward or non-public clinical environments. The presented centralized fleet manager is capable of handling heterogeneous fleets and multi-task robots (see Gerkey and Mataric [GM04]). It mainly consists of a routing engine used for dynamic path planning, a task scheduler solving the task allocation problem, and a controller for executing the resulting work plan. Task allocation is achieved using an Integer Linear Program (ILP) model, which is either solved directly for small problem instances or by means of a sub-optimal genetic algorithm. The optimal solution requires a short computation time of 0.12 seconds for the smallest investigated problem instance (5 robots, 20 locations) and rapidly increases for larger problem sizes (256.58 s for 10 robots and 40 locations, 7201.08 s for 15 robots and 60 locations). Using the genetic algoritm solver, the computation time for larger problem sizes is considerably shorter for larger problem instances (9.56 s for 5 robots and 20 locations, 21.02 s for 10 robots and 40 location, 188.18 s for 15 robots and 60 locations).

Valner et al.

Valner et al. [Val+22] specifically focus on logistical robots transporting samples from intensive care units to laboratories. A real-world hospital in Estonia is used as a reference environment, with its floor plan and environmental characteristics (semi-autonomous doors, moving persons etc.) serving as a foundation for the investigations. For fleet management, the authors use a third-party open-source solution, which is part of the Robot Operating System 2 (ROS 2)[2]. While the centralized fleet manager is capable of handling heterogeneous fleets and multi-task robots, the authors do not explore the merits of these concepts in the context of the application scenario. While the authors initially state to have deployed a heterogeneous robotic fleet to the real-world hospital environment, only a single robot (TIAGo by PAL Robotics) appears to be used for the demonstrations, according to later sections of the paper.

2.2.2 Other Domains

Numerous approaches have been proposed for similar problems in other domains, such as industrial manufacturing [Pfr+14], logistics [BBT01; Liu+14; LVS16; TAV18; Cac+20; CCK20; KP20], traffic management [Xia+13; Bue+17], environmental monitoring [CV10], aviation [SKM14], customer service [Vel+15], emergency services [LK16] and the handling of hazardous tasks [HZC16]. Extensive reviews are provided by Khamis et al. [KHE15] and Seenu et al. [See+20].

[2] https://osrf.github.io/ros2multirobotbook/intro.html

3

Aims and Concept

3.1 Limitations of the State-of-the-art

With respect to the related work summarized in the previous chapter, several limitations of the presented approaches were identified. Due to these limitations, available solutions are not fully applicable to the OR wing, since the unique requirements and characteristics of this environment and the processes taking place within it are not addressed. This mainly pertains to the following aspects:

1. **Context-dependent Prioritization:** In the OR wing context, it is necessary to prioritize pending tasks according to their individual importance and urgency. Especially in situations of high demand, the fleet management system must be able to decide which tasks should be executed first and for which tasks a delay is acceptable. In the most critical cases, the execution of certain tasks may directly affect the well-being of a patient currently undergoing surgery. Clearly, such critical tasks must be marked with a high priority and executed with preference, instead of strictly following a simple first-in-first-out strategy or optimizing only for economic parameters. Since it is not relevant to their use cases, none of the state-of-the-art approaches provide the means for task prioritization or study this aspect and the resulting impact on fleet behavior and performance.
2. **Multi-task Mission Planning:** As further explained in section 4.4.3, robot task allocation and mission planning for the OR wing introduces a novel optimization problem, which can be seen as an extension of the Vehicle Routing Problem. While related problems for other application scenarios within the hospital have been presented in state-of-the-art work, there are key differences, which prevent the application to the OR wing scenario. Among other things, this relates to the presence of alternative pickup locations, which are not covered by the standard formulation of the VRP-variant *Vehicle Routing Problem with Pickups and Deliveries*. Due to the typical repetitive structure of OR wing environments, there are often multiple storage options where a requested material could be picked up by a robot. When assembling a multi-task mission from individual task requests, it is thus reasonable to select storage locations such that the overall performance of the mission is optimized, e.g. with regard to driving durations. While providing a considerable potential benefit, alternative pickup locations are not

considered by the routing or mission assembly strategies of current approaches. Instead, the locations of all points-of-interest relevant to a given task are assumed to be fixed and predetermined at the time the request is made. In the OR wing context, this is expected to lead to unnecessarily long execution durations and thereby cause delays within the surgical workflow.

3. **Response Times:** While most of the presented approaches offer a real-time online fashion of fleet management (as opposed to, e.g., offline shift planning), the achieved computation times for medium to large problem instances are still too long for OR wing requirements, even when using the sub-optimal solving strategies presented in some publications. Since high-priority tasks must be executed immediately (see number 1.), only a very short time-span is available for revising the current fleet schedule and integrating new tasks. Ideally, the reaction time of the system should be comparable to that of a person, which is below 800 milliseconds according to [DD05]. This threshold is confirmed by expert interviews with surgeons that were conducted during the requirements analysis phase of the AURORA research project. The response times achieved for state-of-the-art approaches are at least one order of magnitude higher than desired, with computation times of 20 seconds and higher for medium-sized problem instances [Guz+21; BXD10; BX09].

4. **Proactive Behavior:** In order to cope with the challenging requirements of the surgical domain regarding response times and workload, MSR fleet management should behave proactively by anticipating future user needs and preparing accordingly. For example, materials that are likely to be needed by the surgical team in the immediate future can be proactively fetched ahead of time in order to reduce reaction, driving and execution durations. So far, the merits of proactive fleet management have not yet been considered by any of the available solutions, which are thus limited to merely reacting to outside requests.

Beyond the above domain-specific aspects, methodical weaknesses can be identified. Most of the summarized approaches work with abstract models that are arguably rather far removed from the characteristics of the clinical target environments and workflows. Only few approaches (namely [RKF98; Val+22]) use realistic floor-plans based on actual hospital facilities. For the other approaches, it is therefore unclear, how well the fleet will perform in real-world layouts. Furthermore, none of the approaches conduct studies to observe and record real-world workflows in order to use this data to create realistic task generation models, which can later be used for simulating the workload that is actually imposed on the fleet over time. Thus, it cannot be demonstrated whether the robotic fleet and the fleet management system are capable of handling real-world demand. In any case, it is argued that the OR wing environment, which is the focus of this doctoral thesis, deserves separate consideration both regarding the environmental layout and the workflows taking place within it.

Moreover, the state-of-the-art only provides very limited insights regarding optimal fleet size, fleet composition and robot capabilities such as driving speed and battery duration. The few exceptions, such as [JLK17], focus on environments different from the OR wing. The investigation of such aspects is decisive when it comes to real-world application and necessary for demonstrating that a given MSR fleet is indeed able to handle the occurring workloads.

3.2 Aims of Own Work

As motivated in the Introduction chapter, MSRs are urgently needed as technological means for addressing current challenges within the healthcare system. To that end, the overarching aim of this doctoral thesis is to, for the first time, study the OR wing with respect to MSR fleet integration and demonstrate the feasibility and the value of this proposition. The following goals were defined:

1. **Requirements Analysis:** The characteristics, requirements and risks of the OR wing environment and the workflows taking place within it shall be understood and discussed in the light of MSR fleet integration. Differences to other target environments (within the hospital or other domains) shall be identified.

2. **Fleet Dimensioning for the OR Wing:** It shall be investigated how robotic fleets must be composed and dimensioned in order to perform well within the target environment and add value to OR wing workflows. Firstly, this pertains to the number of robots forming the fleet, which is a trade-off between workforce on the one hand as well as costs (acquisition, maintenance etc.) and space requirements on the other hand. Ideally, the number of robots should be chosen such that the fleet can robustly handle the imposed workload caused by regular OR wing operations. Secondly, the performance of heterogeneous fleet compositions shall be studied and compared to the homogeneous case. These aims shall be achieved by means of simulation, using layout and workflow models closely based on real-world observations.

3. **Fleet Management for the OR Wing:** A fleet management approach shall be developed that specifically takes into account the unique requirements of the OR wing and thus overcomes the limitations of other state-of-the-art approaches with respect to this target application. As described in section 3.1, this relates to the introduction of a task prioritization strategy, the implementation of an adequate multi-task mission planning strategy for the OR wing, the achievement of acceptable computation times and the implementation of proactive behavior. The resulting solution shall be evaluated by means of simulation, again using layout and workflow models based on real-world observations, and shall be compared to the performance levels observed for human-only surgical teams without robotic support.

4. **Surgineering:** With reference to the surgineering paradigm proposed in [Feu+19], it shall be identified at which point the optimization of fleet management and individual robotic design reaches its limits and where it is instead more reasonable to re-think environmental structures and processes, such that MSR technology can be integrated with less effort and its performance can be further increased.

3.3 Concept

As a first major effort, the OR wing shall be studied with respect to MSR fleet integration. This shall be achieved by means of expert interviews with healthcare professionals from the surgical field and by means of real-world observational studies conducted within a German university hospital. In this context, it is of particular interest to record relevant workflows and to identify which type of tasks are suitable for MSR-based execution, where such tasks originate from, how often and in which context they are requested and which points of interests and resources are involved in the execution process. For later reference, the performance of human workers during these tasks shall be documented as well.

As a further prerequisite for simulation, a model of the environment shall be created, which includes the room layout (floor-plan), the location and purpose of points-of-interest (storages, devices, hand-over positions etc.), the definition of pathways and the description of storage contents.

These procedural and environmental models shall become the foundation for implementing a simulation tool for evaluating the performance of a freely configurable robotic fleet acting within the modelled environment and facing the task workload imposed by the recorded workflows. The high-level conceptual structure of such a simulation tool is shown in Figure 3.1.

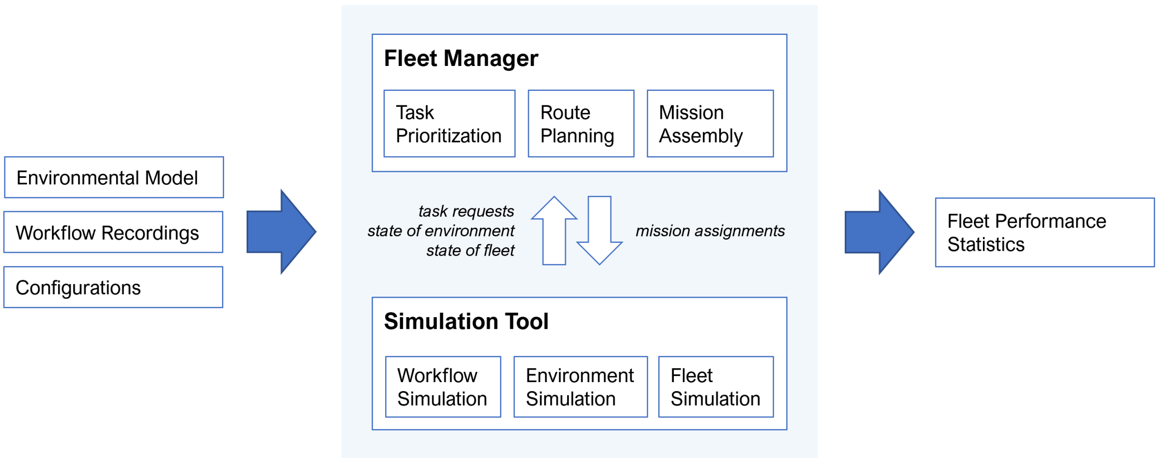

Figure 3.1: A high-level concept of the simulation tool is shown. Based on an environmental model, workflow recordings and a configuration of the simulation scenario (fleet size and composition, storage contents etc.), this tool shall provide means for simulating the state of the environment, the mobile robotic fleet and surgical workflows taking place in multiple operating rooms at the same time. Furthermore, it provides interfaces for communicating with an external fleet manager responsible for task prioritization, route planning and mission assembly.

The simulation software will be used for the purpose of fleet dimensioning and composition as well as for the development, testing and evaluation of fleet management concepts. This shall lead to the implementation of a novel MSR fleet management solution that explicitly considers OR wing requirements and overcomes the limitations of state-of-the-art approaches.

3.4 Novel Contributions

In Table 3.1, contributions of this doctoral thesis and state-of-the-art approaches are summarized for the purpose of comparison and delimitation.

Column 1 indicates which fleet management approaches feature a centralized and which a distributed architecture. The only approach featuring a distributed architecture was proposed by Baalbaki et al in [BX10].

Column 2 describes which approaches are suitable for real-time task allocation, as opposed to offline planning prior to run time. In this regard, Rosetti et al. [RKF98] and Baalbaki et al. (2008) [BLX08] can be categorized as purely offline, while the method proposed by Baalbaki et al. (2009) [BX09] is only feasible for smaller instances, due to computational complexity of the algorithm. All other approaches provide means for real-time task allocation with acceptable computation times for the respective intended applications (albeit not necessarily for the OR wing scenario).

Column 3 indicates which approaches are suitable for managing multi-task robots (i.e. robots that can execute missions consisting of multiple tasks), which applies to all approaches except Baalbaki et al. 2008 [BLX08] and Kumar et al. [KSW18]. Valner et al. [Val+22] use a generic third-party fleet manager that is capable of handling multi-task robots, however, the authors do not utilize this concept and evaluate its merits for the intended use case.

Column 4 indicates which approaches are suitable for composing and managing heterogeneous fleets. In a strict sense, this is only the case for Guzman et al. [Guz+21]. However, the work by Baalbaki et al [BLX08; BX09; BX10] can be considered a special case in this regard since robot bases are identical but can be equipped with different functional modules. Again, the third-party fleet manager used by Valner et al. [Val+22] seems to be capable of handling heterogeneous fleets, but the authors do not explicitly study this aspect.

Column 5 indicates whether the approaches were evaluated using simulation layouts based on real-world hospital environments. This has been done by Rosetti et al. [RKF98] and Valner et al. [Val+22]. In the case of Guzman et al. [Guz+21], it is not clearly stated by the authors whether the used floor-plan was based on a real-world hospital.

Column 6 indicates whether real-world observations of workflows were conducted in order to create a realistic model of task requests over time and whether the fleet performance was evaluated with respect to the resulting workload. This likely applies to Rosetti et al. [RKF98], even though few information is provided by the authors. While Valner et al. have actually deployed a mobile robot to a real-world hospital facility, they do not study the performance of an entire fleet of such systems with respect to workload.

Column 7 indicates which publications are concerned with finding the optimal fleet size and composition for a given scenario. In this regard, Jeon et al. [JLK17] provide results regarding optimal fleet size for their use case. Rosetti et al. [RKF98] compare the performance of three different fleet sizes. The aspect of heterogeneous fleet composition is not studied by any of the approaches.

The remaining columns (8-12) are concerned with aspects that are central to the OR wing use case and, thus, have not yet been addressed by state-of-the-art approaches. As motivated in section 3.1, this pertains to context-dependent task prioritization, multi-task mission planning, acceptable reaction times for OR workflows and proactive fleet management. These gaps are to be closed by the results of this doctoral thesis.

Table 3.1: The contributions of this doctoral thesis (see row "Own Work") are summarized and compared to the contributions of other state-of-the-art approaches. Features that are provided by a given approach are marked in green color. Features that are not provided are marked in red color. Some features may be partially provided or were only studied up to a certain degree, which is indicated with an orange color and further clarified in the footnotes.

	1	2	3	4	5	6	7	8	9	10	11
	Centralized (C) vs. Distributed (D)	**Real-time Task Allocation**	**Multi-task Robots**	**Heterogeneous Fleet**	**Real-world Layouts**	**Real-world Workflows**	**Fleet Dimensioning and Composition**	**Context-dependent Prioritization**	**Multi-Task Mission Planning for OR wing**	**Response Times Acceptable for OR wing**	**Proactive Fleet Management**
Rosetti et al. (1998) [RKF98]	C					a)	b)				
Baalbaki et al. (2008) [BLX08]	C			c)							
Baalbaki et al. (2009) [BX09]	C	d)		c)							
Baalbaki et al. (2010) [BX10]	D			c)							
Jeon et al. (2017) [JLK17]	C						b)				
Kumar et al. (2018) [KSW18]	C										
Guzman et al. (2021) [Guz+21]	C				e)						
Valner et al. (2022) [Val+22]	C		f)	f)							
Own Work	C										

a) The observation of real-world workflows is implied by the authors, however few information is given regarding the methods that were used.

b) Limited results regarding fleet size dimensioning are given.

c) Robot bases are identical, but can be equipped with different modules.

d) The proposed method is only feasible for small problem instances.

e) It is not clearly stated by the authors whether the used floor-plan was based on a real-world hospital.

f) While the utilized generic third-party fleet manager seems to be capable of this, the authors do not use/evaluate these features in their demonstrator.

4

Materials and Methods

In the following, materials and methods used in the course of this doctoral thesis are described. Firstly, the general application scenario is introduced and use cases for MSR within the OR wing are identified (see section 4.1.1). As a basis for data collection and simulation design, a reference intervention, the laparoscopic cholecystectomy, was selected (see section 4.1.2). Section 4.2 describes the recording of workflow data based on real-world observations of the reference intervention. As a methodology for fleet dimensioning and the evaluation of fleet management approaches, the novel simulation framework FleetOR is introduced in section 4.3. This framework was used to develop and evaluate the concepts for managing MSR fleets within the OR wing that are described in section 4.4.

4.1 Application Scenario

While the concepts presented in later sections of this thesis are applicable to many clinical and non-clinical use cases, the focus is placed on the support of workflows taking place in the context of the OR wing. A variety of use cases for supporting surgical procedures and related peripheral processes is identified in Section 4.1.1. While most of these use cases are relevant to arbitrary surgical workflows, a specific procedure was selected for the purpose of data collection and as a reference for testing and evaluation. This type of surgical procedure – the laparoscopic cholecystectomy – is described in section 4.1.2.

4.1.1 Use Cases for MSR within the OR wing

Many types of surgical interventions can be characterized by the following generic workflow: Scrub nurses and circulating nurses are responsible for preparing the operating room for the upcoming surgery. This involves the pre-adjustment of equipment (medical devices, OR table etc.) and the preparation of surgical instruments and materials. Before the patient is brought into the OR, anesthetics are induced by the anesthesiologist. After arrival in the OR, the patient is brought into a suitable position for the surgical intervention to be performed. Surgical drapes and disinfectants are applied before the surgeon performs the first incision. After access to the anatomical target site has been established, the surrounding tissue is dissected to completely expose the pathologically affected structure. The surgeon resects this structure and retrieves it from the patient's body. In a next step, it may be required to reconstruct the anatomy and/or the function of the affected region. Subsequently, the incision is closed and wound dressings are applied. The patient is now brought to the post-anesthesia care unit and further monitored, before being transported back to the ward. Simultaneously, the OR is cleaned and prepared for the next intervention.

The described performance of surgical procedures is an interdisciplinary effort, requiring a diverse team consisting of medical doctors, nurses, technicians and managers. As explained in Figure 4.1, key roles within the OR wing include the surgeon who performs surgical techniques, the surgical assistant who directly supports the surgeon, the scrub nurse who manages the surgical instruments, the circulating nurse who performs non-sterile assistance tasks, and the anesthesiologist who monitors and controls the state of the patient.

Figure 4.1: Different roles and tasks within the OR wing are explained.

According to the findings of the AURORA project and the investigations conducted in the context of this doctoral thesis, promising use cases for MSR within the OR wing mostly revolve around the tasks of the circulating nurse (OR circulator). This role is highly important for the smooth functioning of workflows within the OR, since the other members of the surgical team – such as first surgeon, assistant and scrub nurse – are sterilely washed and dressed, and thus are confined to the sterile zone surrounding the patient. This sterile zone is usually covered by green or blue sheets and includes the OR table, trays and surgical instruments. In case a sterilely dressed member of the surgical team leaves this area, a prescribed hygiene routine must be performed before being allowed to approach the OR table again. This routine includes a precisely defined washing and disinfecting of hands and forearms as well as the donning of a sterile surgical gown. Since this procedure is time-consuming and produces waste, it is usually only performed once before the beginning of the surgical intervention. Intraoperatively, it is only repeated when the temporary absence of a team member cannot be avoided.

At the same time, there is a variety of tasks that must regularly be performed outside of the sterile zone, which motivates the role of an additional non-sterile person – the circulating nurse. Such tasks are often required for proceeding with the intervention and may even be connected to the patient's well-being. Typical examples include:

- **Provision of sterile materials:** Some materials are not prepared preoperatively, since their usage is unlikely or the required variant is not yet clear in advance. Instead, such materials are intraoperatively provided by the circulating nurse on demand. After receiving a request from the surgeon or the scrub nurse, the circulating nurse goes to the storage, takes out the required number of units, documents the usage, returns to the OR table, opens the material and hands it over in a hygienic manner. The hand-over procedure usually follows a well-defined protocol, where the circulating nurse opens the sterile packaging by pulling apart two flaps such that the material inside emerges. Maintaining this configuration, the material is now presented to the receiving person on the sterile side, who can now extract the content without touching any non-sterile parts of the packaging.
- **Adjustment and maintenance of medical devices:** Some medical devices used in the OR are non-sterile (except for specific endeffector parts that may be in patient contact) or possess non-sterile control interfaces and thus cannot be directly controlled by the surgeon for hygienic reasons. Typical examples include the gas insufflator, the laparoscopic light source, the laparoscopic camera control, video monitors, electrosurgical equipment and the OR table control. Following the requests of the surgeon, the circulating nurse is responsible for interacting with these devices.
- **Assistance with sterile dressing:** For hygienic reasons, surgeon, assistant and scrub nurse cannot perform the entire dressing procedure by themselves. They depend on the assistance of the circulating nurse, who closes the backsides of the surgical gowns and helps in tying the ribbons.
- **Phone management:** Since surgeon, assistant and scrub nurse are sterilely washed and dressed, they cannot answer phone calls themselves but depend on external support. Depending on the urgency, the circulating nurse takes the call or holds the phone to the ear of the receiving person.
- **Documentation:** The circulating nurse is responsible for documenting the use of surgical material throughout the surgery. Further documentation tasks include the recording of important points in time such as first incision or wound closure.
- **Specimen and biopsy handling**: The circulating nurse receives surgical specimens from the sterile team and handles them according to the type of intervention. In some cases, the specimen is opened, e.g. to document the number of gall stones during a cholecystectomy. Similarly, the circulating nurse is responsible for receiving tissue samples and forwarding them to the histology laboratory (e.g. by means of tube mail).

Furthermore, the circulating nurse is involved in the preparation and clean-up of surgical interventions. Here, the following types of tasks are commonly performed:

- **Preparation of surgical instruments and materials:** Preoperatively, the circulating nurse assists the scrub nurse in preparing the instrument trays. Since the scrub nurse is already sterilely dressed at this point, the circulating nurse is required to procure sterile materials from the storage and hand them over to the scrub nurse, following the procedure described in the previous.

- **Preparation of medical devices:** The circulating nurse is responsible for moving all required devices into the OR and initializing them according to the requirements of the upcoming surgical intervention.
- **Preparation of patient:** The circulating nurse assists in moving the patient into the OR and positioning him/her on the OR table.
- **Waste disposal:** In the course of a surgical procedure, a large amount of waste is typically generated. The circulating nurse is responsible for counting some types of items, such that it can be assured that nothing remains inside the patient's body, before disposing of them.

While the non-sterile support of the OR circulator is very important for the smooth execution of surgical procedures, it is at the same time often subject to problems, originating from inexperience or miscommunication [Gar+20]. Moreover, a single OR circulator may be responsible for assisting multiple operating rooms, resulting in high workloads and simultaneous task assignments. In turn, this often causes delays during ongoing interventions and leads to frustration of the surgical team, which, in combination with an increasing staff shortage observed for this profession [Mat01; BDO15], further supports the need for automation and technological assistance.

Clearly, not all tasks of the OR circulator are suitable for robotic execution. Some may be better realized by means of software solutions (e.g., phone management, as proposed in [Koh+18]) or are very far removed from the capabilities and safety-levels of current robotic systems (for instance, when considering the positioning of patients on the OR table). However, a large portion of circulating nurse tasks is indeed well-suited for mobile service robotics, even though the technical realization might certainly be challenging for some use cases.

As studied in the framework of the AURORA project, frequent tasks such as the provision of sterile materials and the adjustment of medical devices are of high interest and relevance. Sterile materials must be fetched and handed over both during the preoperative preparation phase and during the intraoperative surgical workflow. In accordance with the findings of the AURORA project, this type of task can indeed be acceptably performed by an MSR. While the adjustment of medical devices may also be solved by means of digital networking, the regulatory hurdles that must be faced are extremely high and the device manufacturers would be required to open digital input interfaces to their products, while adhering to the same standardized protocol. It is highly unlikely that this vision will become reality within the next years, which is why a robotic execution of such tasks can be motivated. The effective and straightforward realization of this proposition has also been confirmed in the course of the AURORA project.

Other logistical tasks, such as the movement of larger objects (instrument sieves, medical devices, OR tables, patient beds, tube mail containers, waste bags etc.) are of high relevance as well. The transportation of similar objects has been studied either in the hospital context – for instance, self-driving patient beds [Wan+16] – or in other application domains – such as factory automation [Hor+11] or package delivery [AIS20] – and thus is expected to be well-suited for MSR-based execution.

While the assistance during sterile dressing is certainly challenging from a technical and hygienic standpoint, it is still a promising use case for MSR technology. While this application has not yet been studied in the OR context, similar approaches were proposed for home care scenarios, such as the work presented by Erickson et al. [Eri+18], which is based on the *PR2* robot platform (Willow Garage, Palo Alto, USA).

4.1.2 Reference Procedure

While the approaches developed in the framework of this doctoral thesis are applicable or extendable to many use cases, a specific surgical procedure was selected for the purpose of data collection and as a reference for testing and evaluation. Such a reference procedure ideally is associated with high case numbers and a standardized workflow, which facilitates data gathering and allows for building a representative data collection based on a realistically attainable number of observed cases. The laparoscopic gallbladder removal (*cholecystectomy*) is a prime example of such a surgical procedure, which is why it was selected as reference intervention.

The gallbladder is a pear-shaped hollow organ that stores and concentrates bile produced in the liver before it is released into the duodenum to aid in fat digestion [ZT10]. The organ is located on the underside of the liver and is connected to the *ductus hepaticus communis* (arriving from the liver) via the *cystic duct* [SS12]. Both ducts unite to form the ductus choledus, which enters into the duodenum [SS12]. Blood supply to the gallbladder is established via the *arteria cystica* [ZT10].

Diseases of the gallbladder may necessitate its removal by means of surgery. One common pathology is the symptomatic cholecystolithiasis, where patients suffer from pain caused by gallstones that have formed within the gallbladder [SS12]. This may be accompanied by complications such as inflammation or perforation of the gallbladder [SS12]. In practice, minimally invasive cholecystectomy performed laparoscopically has largely prevailed over the conventional open variant [Car14]. The procedure is widely performed and thus associated with high case numbers: in Germany alone, approximately 175,000 cholecystectomies are performed each year [IQT21].

The surgical workflow is highly linear and standardized, with a comparatively low number of steps involved [Twi+17], as summarized in the following [SS12; Car14; Sta+16]: CO2 gas is first insufflated into the patient's abdomen such that trocars can be placed. The laparoscope – a rigid straight variant of the endoscope – and other required surgical instruments, such as suction-irrigation device, grasping forceps, etc., are introduced through the trocars. The gallbladder tissue in the area of the so-called *Calot triangle* is now dissected to reveal the arteria cystica and the ductus cysticus. Each vessel is sealed with the aid of a clip applicator before being cut. This is followed by dissecting the gallbladder and separating it from its liverbed using the electrocoagulator. The gallbladder can now be extracted from the abdominal cavity with the aid of an extraction bag. After a final check, the trocars are removed and the incisions are sutured.

While the cholecystectomy was used as a model procedure throughout this thesis, it must again be emphasized that the methods and concepts presented and evaluated in the following are not limited to this type of intervention. Especially more complex interventions are expected to benefit greatly from robotic support, since their workflows are less predictable and more on-demand tasks must typically be carried out. Such types of interventions are, however, less suitable as model procedures due to their lower case numbers and higher workflow variance. On the other hand, if the benefits of fleet dimensioning and management can be shown for highly linear and standardized intervention types, this can be seen as an even stronger support of the value of this technology.

4.2 Data Collection

As a prerequisite for simulation, real-world data was collected regarding the surgical workflow of the reference intervention (see section 4.2.1). Herein, the focus was placed on observing the duties of the circulating nurse, as motivated in the previous. For building a simulation model of the surgical environment, a real-world operating room wing was studied with regard to layout and points of interest, as described in section 4.2.2.

4.2.1 Workflow

To gain a better understanding of OR circulator tasks and the workflow they are embedded in, an observational real-world study was conducted, which later served as a foundation for workflow simulation and as a reference for evaluation. By recording the start and end times of executed tasks, a realistic view on the workload of the OR circulator during the pre-, intra- and postoperative phases was gained. This data was later used to impose realistic workloads on the simulated MSRs and thereby study the performance of different fleet configurations. The recorded data was also used for directly comparing the fleet performance to the status quo performance achieved by human OR circulators without robotic support. In total, the perioperative workflows of 20 laparoscopic cholecystectomies were observed, while recording the following information:

- **Task type:** The type of task performed by the OR circulator (see table 4.1 for a description of all considered task types).
- **Time of request:** The time at which the request has been made by a member of the surgical team.
- **Start time:** The time at which the OR circulator has started to execute the task.
- **End time:** The time at which the OR circulator has completed the task.
- **Parameters:** Task-individual parameter values depending on the task type (see table 4.1 for details).
- **Emergency level:** The urgency of the task, as defined in [Ber+20] and described in section 4.4.2.
- **Command level:** The position within the clinical command hierarchy of the surgical team member making the request, as defined in [Ber+20] and described in section 4.4.2.
- **Perioperative events (timestamps):** Coarse-granular events marking important steps within the patient's treatment workflow (e.g. begin of pre-operative OR preparation,

arrival of patient within the OR, begin of preoperative patient positioning, first incision, wound closure, departure of patient, begin of cleanup etc.).

- **Intraoperative phases (timestamps):** Fine-granular phases marking distinct parts of the intraoperative workflow between first incision and wound closure, based on the definitions in [Twi+17].

Furthermore, it was recorded whether an observed task was executed during the preoperative preparation phase, during the actual intervention or during the post-operative cleanup phase. In some cases, multiple requests were communicated to the OR circulator at the same time, which complicated the differentiation of precise start and end times for each single request. Therefore, only the combined start and end times were recorded for such batches of requests.

Table 4.1: A summary of the recorded OR circulator task types is given. Each task type is described and the task-specific parameters are listed.

Task Type	Description	Parameters
C1	Transportation of sterile goods (e.g. from storage to OR table)	article number, amount, start location, destination
C2	Transportation of a heavy load (e.g. containers, devices)	article number, amount, start location, destination
C3	Intraoperative disposal of waste	start location, destination
C4	Postoperative counting and disposal of waste	start location, destination
C5	Transportation of specimen to preparation table (located within the same OR)	start location, destination
C6	Delivery of lab samples to the OR wing's tube mail station	sample type (swab, frozen section procedure, compound), start location, destination (microbiology, pathology, histology)
C7	Adjustment of medical device	device type, device location, target parameters
C8	Assistance during the sterile clothing procedure	operating room, type of clothing (e.g. full, gloves only)
C9	Management of incoming call	
C10	Management of outgoing call	

For each intervention, the recorded data described above was added to a joint database and augmented with further meta information, such as surgeon ID (anonymized), assistant ID (anonymized), OPS codes, ICD codes and operating room ID. In this form, the recorded data can be accessed and processed by the workflow module of the FleetOR simulation environment (see section 4.3.4).

4.2.2 Environment

In order to build an environmental model for simulation purposes, a real-world OR wing at a German university hospital was studied. The layout of the facilities was extracted from architectural floor plans depicting the locations of operating rooms, storage rooms, corridors and other spaces. This was augmented by further information gained from real-world observations, including the locations and contents of storage cabinets, the location of medical devices and the location of points of interest (POIs), such as tube mail stations and disposals. For each operating room, a generic surgical layout was defined with realistic positions for the handover of surgical material and for the use of different medical devices.

4.3 Simulation

Since MSR technology for the OR wing is still in its infancy and real-world fleets are not yet available, simulation is the means of choice for demonstrating the benefits and identifying the limitations of this technology. To that end, a novel framework for simulating robotic fleets within the OR wing, first introduced in [Ber+22] and referred to as *FleetOR*, was implemented for conducting the studies presented in the Results chapter. In the following, the requirements of this simulation framework are described in section 4.3.1, followed by the chosen software architecture (section 4.3.2) and its main modules for simulating the OR wing environment (section 4.3.3), the surgical workflow (section 4.3.4) and the robotic fleet (section 4.3.5).

4.3.1 Requirements

Depending on the insights to be gained, the simulation of robotic systems may take various shapes and forms. One might be interested in visualizing the movements of a robotic manipulator, testing collision avoidance of mobile bases or generating virtual sensor data. In the context of the work presented in this doctoral thesis, a *global* point of view regarding robotic fleet behavior is adopted, targeting high-level performance parameters. These results are intended for solving high-level objectives, such as finding optimal fleet sizes, fleet compositions and storage placements. Consequently, the considered entities (facilities, robots, workflows) are modeled and simulated on a high level, neglecting some fine-granular aspects. While these limitations will be indicated throughout the following sections, their influence is assumed to be small compared to higher-level parameters.

From a functional standpoint, the simulation environment must provide the following features to be suitable for MSR fleet simulation within the OR wing:

- Model the environment of the OR wing and update its state over time
- Simulate the execution of surgical workflows
- Simulate the behavior of the robotic fleet members
- Track the state of all relevant entities over time
- Calculate performance parameters (task execution durations, total intervention durations etc.)

- Offer interfaces for flexibly varying key parameters of the simulation (fleet size and composition, robot capabilities, OR schedules, storage locations and contents etc.)

While a detailed graphical 3D visualization – as provided by simulation software such as *Unity*[1] (Microsoft, Redmond, Washington, USA) – may be helpful for debugging and demonstration purposes, it is not required for generating the numerical results presented in this doctoral thesis and would introduce avoidable computational overhead. Nonetheless, some condensed form of visual representation is desirable to monitor the fleet behavior and gain insights regarding the impact of certain changes to the simulation parameters. Thus, it is desirable to be able to execute and visualize the simulation progress in wall-clock time.

Multiple simulation frameworks exist that can be utilized for the purpose of simulating robotic fleets, such as *Gazebo*[2] (Open Source Robotics Foundation, Mountain View, California, USA), *SimPy*[3], *Plant Simulation*[4] (Siemens Digital Industries Software, Plano, Texas, USA) or the aforementioned Unity. While these frameworks use methodologically very different approaches, none of these proved to be a good fit for the work presented herein. Gazebo and Unity place a strong focus on low-level 3D representation, which is not relevant for the purposes of this doctoral thesis. Simpy is limited when it comes to real-time simulation. Industrial software suites, such as Plant Simulation, are highly capable yet rather far removed from the OR wing use case to be effectively used. For these reasons, a dedicated simulation framework was implemented from scratch. While this was associated with a major development effort, it has facilitated the creation of a tailor-fitted solution where all components and parameters can be adjusted and extended freely.

4.3.2 Software Architecture

As shown in Figure 4.2, the developed simulation framework FleetOR consists of the following main modules:

- **Facility Simulation**: Based on a model of the OR wing environment, the states of storages, robot bases, points of interest, paths and zones are simulated over time. The module provides an interface to the fleet manager (see section 4.4) to make the current state of the environment available to the mission planning process.
- **Workflow Simulation**: Using this module, concurrent surgical workflows taking place in multiple operating rooms can be simulated based on recorded intervention data (see section 4.2). Recorded task requests originating from the surgical team are sent to the fleet manager for mission planning.
- **Fleet Simulation**: This module spawns and manages robot simulator instances for each fleet member. Each robot provides an interface for receiving allocated tasks and other instructions from the fleet manager.
- **Database**: The database contains recorded intervention data used by the workflow simulation module as well as further information used for visualization purposes.

[1]https://visualstudio.microsoft.com/vs/unity-tools/
[2]https://gazebosim.org/home
[3]https://simpy.readthedocs.io/en/latest/
[4]https://www.plm.automation.siemens.com/global/en/

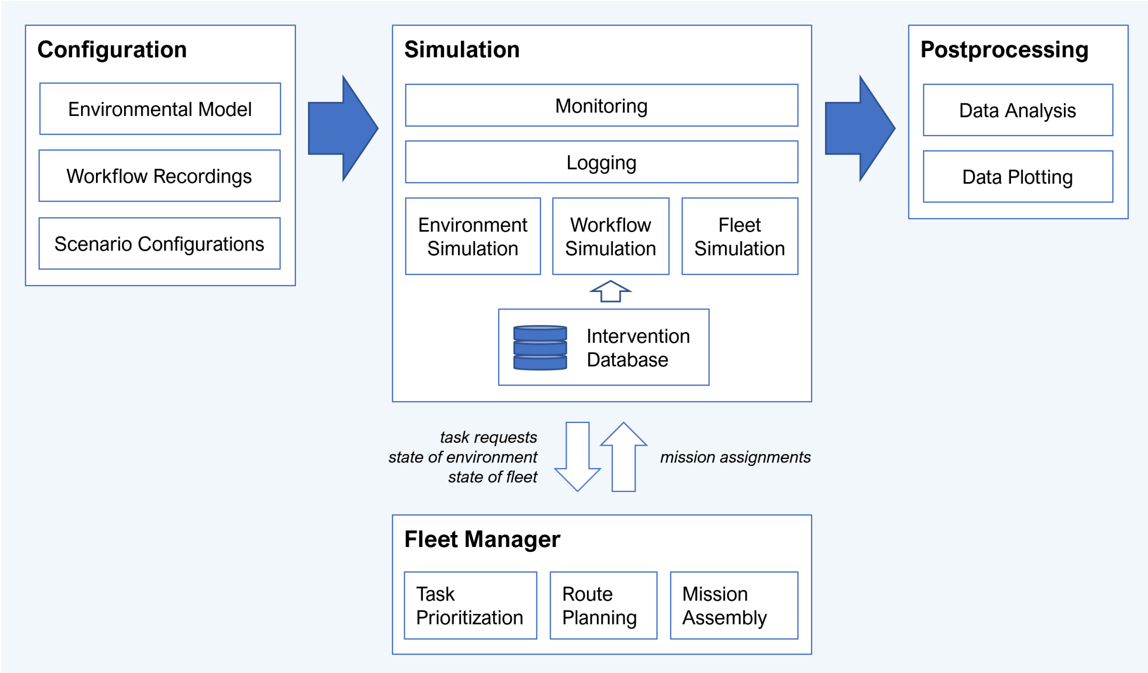

Figure 4.2: The software architecture of the developed simulation framework FleetOR is shown. Based on an environmental model, workflow recordings and scenario configurations, the simulation is executed by the main modules dedicated to the OR wing environment, the surgical workflows and the robotic fleet. The workflow simulation module queries the recorded workflows of the configured interventions from the intervention database. The logging module records the state of simulated entities over time for later analysis and plotting during the postprocessing phase. The monitoring module is used for visualization and debugging purposes. The simulation interacts with the fleet manager by sending task requests and states of the environment and the robotic fleet. In return, mission assignments for the simulated robots are received.

- **Monitoring**: This module provides a dashboard summarizing the current progress of the simulation as well as the current state of all relevant entities. Furthermore, a 2D top-down visualization of the OR wing environment displays robot positions, paths and points of interest such as storages and robot bases. This module is mainly used for demonstration and debugging purposes.
- **Logging**: This module logs the state of relevant simulation entities (robots, tasks, interventions) over time for later evaluation.
- **Data Analysis**: This module parses the log files and generates a summary of the simulation results. For statistical analyses, the Python package NumPy was used[5].
- **Data Plotting**: This module plots the simulation results based on the output of the data analysis module.

For reasons of performance, all core simulation modules (facility simulation, workflow simulation, fleet simulation, logging) were implemented in C++11. Other components responsible for creating and managing simulation instances, as well as those used for post-simulation analysis and plotting, were implemented in Python 3. The monitoring module

[5]https://numpy.org/

was implemented as a web application (AngularJS[6], Bootstrap 4[7]) and connects to the core simulation modules using Rosbridge's roslibjs[8,9].

The entire simulation framework is encapsulated using Docker[10] images, which allows for creating and running multiple simulation instances concurrently. Using a global parameter, the simulation speed can be adjusted, which allows for simulating long-running processes within a short amount of time.

All simulations were run on a 3.5 GHz 12-core Intel Core i7-5930K CPU with 32 GB of RAM and 32 GB of swap space.

4.3.3 Facility Simulation

Fleet management is heavily depending on the nature and the current state of the environment in which the robotic fleet is acting. This includes information such as the positions of POIs and the current inventories of storages. Based on that, decisions are made by the fleet manager, e.g., which robot, given its current position, can fulfill a certain task with minimal driving duration.

To simulate the interaction of fleet manager and OR wing environment, a facility simulation module was developed, closely modelling an existing OR wing at a German university hospital. For that, the location and inventory of storages as well as the location of POIs (devices, disposals, tube mails etc.) were modeled based on a real-world survey (see section 4.2.2). Robot bases and paths were added based on a floor plan of the facilities (see Figure 4.3).

Storages

Storages are central elements of the environmental model, since many task types that are suitable for MSRs within the OR wing involve interactions with some kind of storage facility. Based on examinations of the real-world OR wing mentioned in the previous, the locations and inventories of all storages used for laparoscopic surgical interventions were documented and made available to the simulation software using a JSON-based description format. This description file is used to initialize the simulated storages with the specified article numbers and amounts. While the facility simulation module is able to track changes in stock amounts over time, an infinite supply of the specified types of articles was assumed for most of the simulations conducted herein. Effectively, this expects the restocking of storages to be guaranteed by peripheral processes that are not part of the simulation.

Points of Interest

Beyond storages, there are certain POIs relevant to the task types considered herein. In contrast to storages, POIs are not associated with inventories, but merely mark important locations, such as the following:

[6] https://angularjs.org/
[7] https://getbootstrap.com/
[8] https://wiki.ros.org/roslibjs
[9] https://wiki.ros.org/rosbridge_suite
[10] https://www.docker.com/

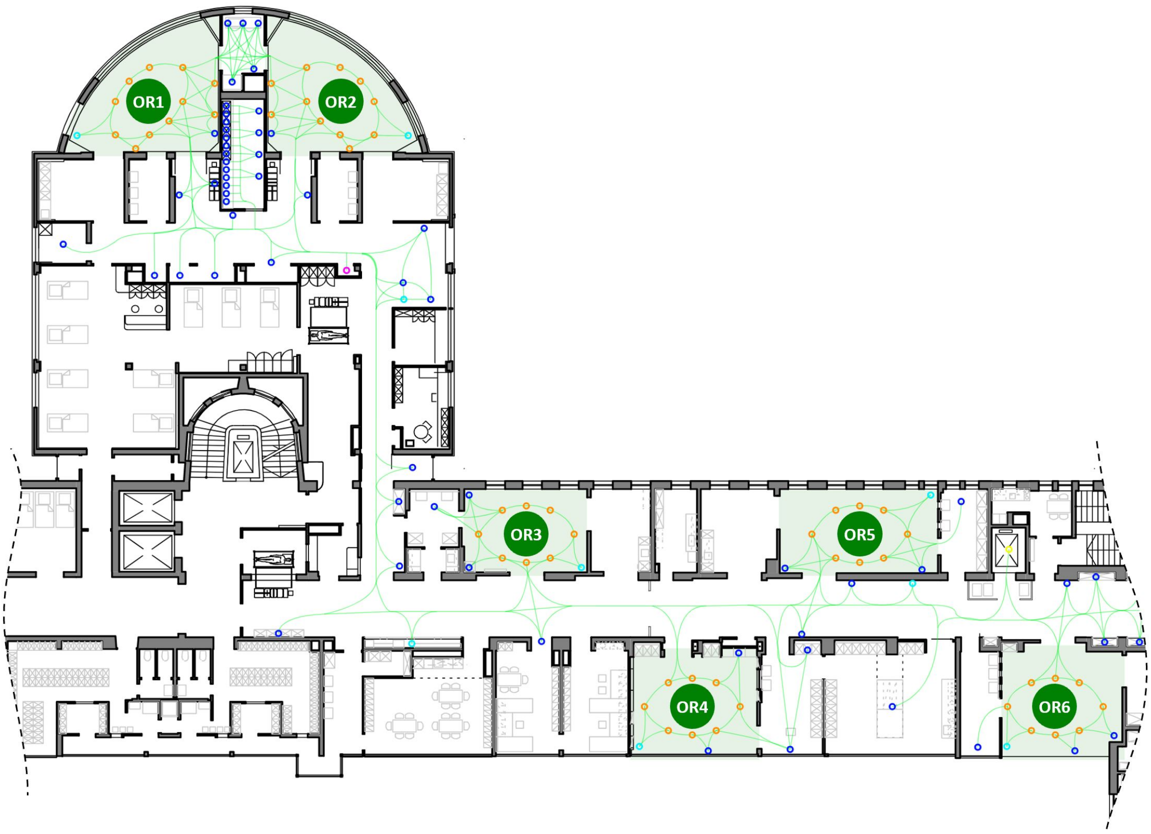

Figure 4.3: A visual representation of the model used for simulating the OR wing environment is shown. POIs are depicted as orange-colored circles, storages as blue-colored circles, robot bases as cyan-colored circles and robot paths as green-colored trajectories. Operating rooms are shaded in light green. (Figure from [Ber+22])

- **OR table**: This POI type defines the designated position for handing over sterile material to the surgical team. While this position may vary depending on the type of intervention and the required OR layout, a fixed handover position was defined for each OR in order to simplify the simulation setup.
- **Device**: This POI type defines the location of a medical or non-medical device (e.g., insufflation device, electrocauterizer, OR table control, OR light control, room light control, suction/irrigation device, monitor control, laparoscopic video processor, cold-light source, SAP computer). Each OR is equipped with a dedicated set of these devices.
- **Disposal**: Waste bins are located within each operating room for intra-operative use. Larger containers are available in the OR wing periphery for post-operative disposal of waste bags.
- **Tube mail**: In the OR wing context, tube mails are used in some hospitals to send tissue samples to the laboratory for further examination. In accordance with the observed real-world OR wing, two tube mail stations were included in the model.

Paths

Determining optimal robot routes is a central part of fleet management, since one objective during mission planning is the minimization of driving time (see section 4.4.3). To calculate and compare travel distances of alternative mission plans, the fleet manager requires a path planning module, which is able to estimate the robot's path, given desired start, intermediate, and end locations. Since the actual path later taken by the robot is subject to uncertainties (e.g. due to the avoidance of unexpected obstacles), it will most likely differ from the "ideal" path assumed at the time of mission planning. However, on average, it is to be expected that these unpredictable influences are the same for all paths, and therefore, relying on an idealized path network is justifiable when making relative comparisons between path lengths.

Based on the floor plan of the studied real-world OR wing, a path network was defined for the purpose of fleet management (see Figure 4.3). The vector graphics editor *Inkscape*[11] was used for drawing the path trajectories using linear segments and cubic bezier segments. The path descriptions were then extracted from the vector graphics format and converted to a JSON-based representation, which serves as an input for the path planning module (see section 4.4.3). In total, the modeled path network consists of 414 path segments connecting 172 POIs.

The path network is also used by the fleet simulation module (see section 4.3.5) for simulating the movement of robots. Robots are assumed to exactly follow these paths when moving between locations of interest. While this is a difference to the expected real-world behavior – where driving times may be elongated due to obstacle avoidance – this aspect is neglected for the higher-level considerations made in the context of this doctoral thesis.

Zones

The facility simulation module provides the possibility to divide the OR wing environment into zones with differing properties. In particular, this was used to define zones of different robot driving speeds, with a lower speed limit for operating rooms and a higher speed limit for corridors and peripheral rooms. The motivation for considering such constraints and the obtained results are described in section 5.3.

When configuring the facility simulation module, speed zones of rectangular shape can be defined using a JSON-based format. The speed limit is applied to path segments with both end points located within the zone. The fleet simulator (see section 4.3.5) can be set up to adhere to these constraints when simulating the movement of robots through the environment.

Robot Bases

MSRs are usually battery-powered and thus have a limited operating time before a recharge is required. To avoid manual interactions to start and stop the charging process, robot bases can be installed within the environment (consider, e.g., home stations of robotic vacuum cleaners). Such bases allow for the automated docking and subsequent recharging of robots. For instance, an *etaLINK 3000* (Wiferion GmbH, Freiburg, Germany) charging station is used for the AURORA robot.

Since the positions of such bases are relevant to fleet management, they can be made available to the facility simulation module using a JSON-based description format. The fleet manager receives this information from the facility simulator and uses it to determine the closest base in situations where a given robot needs a recharge or is currently unoccupied and must be parked at a suitable position.

In the context of the simulations presented in chapter 5, a base may allow for the docking of more than one robot, which is why each base maintains an internal list of currently coupled

[11] https://inkscape.org/about/

robots. The facility simulator allows for freely adjusting the type and number of robots a base can host at the same time. Furthermore, the base provides means for querying available docking slots and for checking in and checking out fleet members.

4.3.4 Workflow Simulation

The workflow module is used for simulating multiple surgical interventions taking place in different ORs concurrently. To achieve that, the module executes an intervention schedule, which can be assembled from recorded cases, and spawns task requests according to the recorded workflows.

Intervention Schedules

Recorded cases (see section 4.2) can be freely combined and assigned to ORs in order to define an intervention schedule that is to be executed by the workflow simulation module. Using the IDs of recorded cases, it is possible to assemble a day's intervention schedule for each OR using a simple JSON-based description format:

```
{
    "schedule": {
        "OR_ID_2": [
            "intervention_ID_7",
            "intervention_ID_1",
            "intervention_ID_6",
            ...
        ],
        "OR_ID_4": [
            "intervention_ID_2",
            ...
        ],
        ...
    }
}
```

Task Requests

Task requests express a complete, clearly defined and delimited assignment from the perspective of the user. As such, they represent a self-contained process, which may subsequently be decomposed into atomic tasks and reassembled to a multi-task mission by the fleet manager (see section 4.4.3). Task requests may originate from different sources, such as voice interaction systems, touch-enabled GUIs or other types of user interfaces. Workflow support algorithms may be a further source of task requests, e.g. when predicting future user needs or when preparing for an upcoming intervention.

Each task request must provide the following information:

- **ID:** Unique ID of task request.
- **Type:** Type of task request, as defined in Table 4.2.

- **Parameters:** JSON-formatted string providing type-dependent parameters required for mission planning and execution (see Table 4.2).
- **Command Level:** Rank of the staff member making the request, as defined in [Ber+20] and described in section 4.4.2.
- **Emergency Level:** Urgency of the request being made, as defined in [Ber+20] and described in section 4.4.2.
- **OR ID:** ID of the operating room from which the task request is originating.
- **Timestamp:** Time at which the request was entered into the system.

For the purposes of the simulations presented in chapter 5, a high-level categorization of task types was developed (see Table 4.2 and cf. [Ber+22]), which reflects the insights presented in section 4.1.1. Robots communicate their ability to execute certain task types by listing them within their capability description (see section 4.3.5), which is broadcasted to the fleet manager.

Table 4.2: The different task types considered within the conducted simulations are summarized. Each task type is described and the task-specific parameters are listed.

Task Type	Description	Parameters
T1	Transportation of a sterilely packaged material (e.g. from storage to OR table)	article number, amount, start location, destination
T2	Transportation of a heavy load (e.g. containers, devices)	article number, amount, start location, destination
T3	Intra- or postoperative disposal of waste	start location, destination
T4	Delivery of lab samples to one of the OR wing's tube mail stations	sample type, start location, destination
T5	Adjustment of medical device	device type, device location, target parameters
T6	Assistance during the sterile clothing procedure	operating room

Workflow Execution

A dedicated OR simulator instance is created for each OR and initialized with the desired intervention IDs according to the configured intervention schedule. Each instance queries the recorded data of its assigned interventions from the database (see section 4.3.2). This includes the following information:

- **ID of first surgeon** (anonymized)
- **ID of assistant** (anonymized)
- **Type of intervention**

- **OPS Code**[12]
- **ICD Code**[13]
- **Tasks:** Recorded actions of the OR circulator, as described in section 4.2.1
- **Events:** Timestamps of important workflow events (e.g. first incision, wound closure), which are routinely documented within the clinical information system
- **Phases:** Timestamps of high-level intraoperative phases, which are not routinely documented, but were recorded additionally

While some of this information (staff IDs, intervention type, OPS codes, ICD codes, phases) is used for visualization or debugging purposes, the task and event recordings are used for the actual workflow simulation. For that, the OR simulator processes the tasks and events in a consecutive order.

The first event marks the beginning of preoperative preparations, which is associated with a list of preliminary tasks, mostly revolving around the collection and preparation of sterile materials, tools and devices. In contrast to intraoperative tasks, which are performed on demand according to ad hoc needs of the surgical team, these preparation tasks are not required to be executed in a certain order. Thus, the OR simulator sends an unordered set of task request messages to the fleet manager, which takes care of task prioritization and mission planning (see section 4.4). It is important to mention that there is usually not a clear deadline for the completion of these preparation tasks, since the beginning of subsequent workflow steps of the recorded intervention may have depended on other factors, not only the completion of OR preparations (e.g. delays during patient transfer, late arrival of OR team members). Therefore, in case task execution is finished before the next recorded workflow event (marking the arrival of the patient at the OR), the workflow simulation does not move on and instead waits for the remaining time span. This mechanism was implemented to achieve a better comparability to the human-only benchmark. However, in cases where robotic task execution takes longer than expected and the beginning of the next workflow event would be overrun, the entire subsequent workflow is moved backward, effectively leading to an elongation of the simulated intervention, compared to the underlying recording of the real-life intervention.

Starting with the workflow event marking first incision (German: "Schnitt"), the simulation is switched over to an "on demand"-style execution, since single tasks (or small sets of tasks) may now be requested by the surgical team at arbitrary times. Here, the simulator simply waits for the next recorded task, sends an according request to the fleet manager, and waits for the simulated robotic fleet to finish the task execution, before again waiting until it is time to request the next recorded task.

[12] *Operationen- und Prozedurenschlüssel*: A coding system published by the German Ministry of Health, which is used for classifying the type of conducted interventions/procedures (also see https://www.bfarm.de/DE/Kodiersysteme/Klassifikationen/OPS-ICHI/OPS/_node.html)

[13] International Statistical Classification of Diseases and Related Health Problems (also see https://www.bfarm.de/DE/Kodiersysteme/Klassifikationen/ICD/ICD-10-GM/_node.html)

4.3.5 Fleet Simulation

The fleet simulation module allows for configuring and simulating the behavior of MSRs within the model OR wing. The fleet can be composed of an arbitrary number of members, which may differ in capabilities and other key performance properties. Each member provides an interface for receiving task assignments, which are subsequently executed by the simulated robot. This encompasses typical robotic activities, such as driving, manipulating, charging or waiting.

Robot Description

As described in section 4.1, the OR wing provides various use cases for MSRs, mainly revolving around tasks of the circulating nurse and peripheral logistics. While, from the perspective of space availability and flexibility, it would be preferable to have multi-purpose robots capable of executing several different types of tasks, such systems may be more complex, costly, error-prone and cumbersome to develop. On the other hand, spreading capabilities across the fleet by incorporating different robot types leads to an increased fleet size required for providing the same assemble of capabilities. This introduces higher acquisition and maintenance costs as well as increased space demands.

In order to study such aspects of fleet composition, the fleet simulation module provides means for defining different types of robots. For each fleet member, the following properties can be configured:

- Driving velocity
- Battery capacity
- Battery charge rate
- Battery drain rates for different activities
- Capabilities
- Execution durations for different activities
- Initial content of onboard storage (robot inventory)

A robot's capabilities can be defined on different levels of granularity, ranging from atomic low-level actions to the high-level collection of services it offers to the user. Making the capabilities of robotic fleet members available to the fleet management system is an essential prerequisite for task allocation, since it allows for dynamically matching requests to the currently most capable agent. This process can become almost arbitrarily complex, depending on the parameters considered (e.g. load volumes, workspace constraints), and introduces the need for describing robotic capabilities in a structured and machine-readable fashion. Developing capability (and corresponding task) description formats has been the subject of a wide range of scientific work [ZAF16]. For the purposes of this doctoral thesis, a code-based approach was chosen, where a unique code is assigned to each possible type of task request (see section 4.3.4). As part of its publicly broadcasted member description, each robot exposes the subset of these codes it is capable of executing. Consequently, the robot is only considered for these task types by the fleet manager. All defined robot types are described in Table 4.3 (also cf. [Ber+22]).

Table 4.3: All robot types studied within this doctoral thesis are listed. For each type, a unique identifier is given, along with a list of task types (see Table 4.2) that robots of this kind are capable of executing.

Robot Type	Capabilities	Description
R1	T1, T3, T4	Transport robot for light loads
R2	T2	Transport robot for heavy loads
R3	T5	Robot for adjusting medical devices
R4	T6	Robot for assisting the sterile clothing procedure
R5	all (T1-T6)	All-rounder robot

Robot Status

Next to the description of its capabilities, a robot's status is another crucial information for fleet management. Since, in the context of modern OR wings, seamless networking can be taken as a given, it is assumed in the following chapters that each robot is able to continuously make its status available to the fleet manager. This may encompass information such as:

- **Availability:** This entity describes the current readiness of the robot to receive and execute a new task. The robot becomes unavailable as soon as a new task is accepted or in case it requires a recharge. After task execution is finished or a sufficient battery level is reached, it becomes available again.
- **Position:** This entity describes the current position of the robot within the OR wing environment as x-y-coordinates. If the robot is currently at a named location (storage, base etc.), the ID of this POI is provided as well.
- **Current activity:** This entity provides information on the task that is currently executed by the robot. This is mainly used for monitoring purposes.
- **Battery level:** This entity is used by the fleet manager to send the robot back to one of the bases should the battery level drop below a certain (adjustable) threshold.

Task Execution

Upon receiving a new task assignment described by an atomic task (see section 4.4.3), the robot's execution routine is triggered. In case the robot is currently docked at one of the bases, it is checked out before driving to the target location of the atomic task. For that, the shortest path from the robot's current position to the target location is queried from the path planning module (see section 4.3.3, subsection *Paths*). Based on the path, the driving duration is calculated, considering the robot's driving speed and the speed limits introduced by special zones within the environment (see section 4.3.3, subsection *Zones*).

In case the monitoring option is activated (see section 4.3.6), a movement along the path trajectory is simulated by continuously publishing position updates in an adjustable interval. While this method generates the intermediate positions necessary for visualizing the robot's movement within the map of the OR wing environment, it is associated with additional overhead, which may skew the simulation results for high simulation speeds. In case the monitoring is deactivated, the robot simulator simply waits for the calculated driving duration and updates its position only once to the coordinates of the target location. This approach generates less overhead and provides more accurate simulation results for high simulation speeds. It is therefore used whenever monitoring is not required or when final simulation results need to be generated.

After the target location has been reached, the action part of the atomic task is executed. For that, the robot may need to enter a queue, before it is allowed to access the POI (e.g. the storage), as described in subsection *Queuing*. After access has been granted, the actual performance activity is simulated, which commonly involves some kind of end effector manipulation at the POI location. Again, this is achieved by waiting for a certain duration, which may vary among task and robot types. For example, a simple button press may take less time than a complex hand-over interaction with a member of the surgical team.

After both the driving and the action parts of the atomic task are finished, the fleet manager is notified regarding the successful completion of the assignment. The robot's availability flag is now set back to true and a new task assignment or an order to go back to base is awaited.

It is important to emphasize that the execution of an atomic task is done in a non-preemptive manner, meaning that it cannot be paused and resumed for the execution of other tasks with higher priority. While a preemptive behavior may be desirable in some exceptional (e.g. emergency-related) situations in order to achieve shorter reaction times, the reasonable preemption of tasks is complex and may even introduce additional overhead during task execution: Consider, for example, a robot opening and handing over a requested sterile material to the scrub nurse. Let's suppose that after the material has been opened, a higher-priority task is received from the fleet manager and the robot is asked to preempt the execution of its current task. The resolution of this situation is not straightforward: Should the robot place the opened material into its magazine for a later handover? Or should the material be thrown away to avoid contamination during the execution of the higher-priority task? In the latter case, the material must be brought to the trash and a new one must be fetched, which increases the workload of the robot. To avoid the resolution of such situations, atomic tasks are considered to be non-preemptive in the context of the simulations presented in chapter 5. Since atomic tasks already represent a very fine-granular segmentation of robotic workflow, this limitation is certainly acceptable for high-level simulations.

Battery Simulation

Since mobile robots are usually designed for operating in an untethered configuration, batteries are required for power supply. This circumstance effectively limits the operational time of the robot and therefore is an important constraint to be considered during fleet management. From a global perspective, the combined battery constraints of the individual fleet members limit the workforce of the entire fleet and therefore have a direct impact on task execution durations and overall fleet performance. To investigate the influence of this important aspect in the context of surgical workflows, a battery simulation was implemented as part of the robot simulator.

During the execution of tasks, the following equation is used to drain the robot's battery, with drain rate d, the battery level b_i at the current time step t_i, the battery level b_{i-1} at the previous time step t_{i-1}:

$$b_i = b_{i-1} - d(t_i - t_{i-1})$$

To model the specific power consumption associated with various phases of task execution, specific drain rates are used for driving (d_d), performing (d_p) and waiting (d_w). Should a robot's battery level drop below a predefined threshold b_{low}, it autonomously returns to its base for recharging. Recharging continues at least until the battery level reaches b_{high}. In the meantime, the robot remains unavailable for task execution. Upon completing a mission with no further pending tasks, the robot is directed to the nearest unoccupied base for recharging, irrespective of its current battery level. The recharging process is modeled accordingly, incorporating a variable charging rate (c) that may differ between robots:

$$b_i = b_{i-1} + c(t_i - t_{i-1})$$

Queuing

During task execution, robots may be required to use shared resources within the OR wing environment, which means that access to these resources must be coordinated. It is assumed that only one robot can interact with a POI – such as a storage, device, disposal etc. – at the same time. Therefore, a coordinating mechanism is required that enqueues newly arriving robots and sequentially grants access. For the sake of the simulations presented in chapter 5, the following strategy was implemented: Incoming robots navigate to the POI and request access. For the trivial case that the POI is currently vacant, access is granted immediately. In case the POI is already in use by another fleet member, the robot's ID is added to an internal queue representation. The robot is now required to wait until all previous fleet members have left the queue, before, finally, access is granted.

Different queuing strategies may be implemented. While the most trivial of such is a simple first-in-first-out (FIFO) approach, further contextual knowledge may be used to implement a more sophisticated access control. For example, incoming robots can be prioritized according to the urgency of their current mission. As during mission planning, such a prioritization can be based on the task's emergency level, command level or idle time (see section 4.4.2).

4.3.6 Monitoring

While the monitoring module of FleetOR is highly useful for debugging purposes as well as for visualizing the processes taking place within other modules of the simulation, it is not a mandatory component for running simulations. Moreover, it introduces additional I/O overhead, which may lead to a distortion of simulation results, especially for higher simulation speeds. Therefore, the monitoring module was implemented as a separate component that can be flexibly attached and detached from the main simulation.

To benefit ease of use and avoid the necessity of installation, the module was implemented as a web-based application. For the frontend, common web-frameworks were used, in particular *Bootstrap 4*[14] and *AngularJS*[15]. The web application is served from an *Apache*

[14] https://getbootstrap.com/
[15] https://angularjs.org/

HTTP server running inside a *Docker container*[16]. The frontend connects to the core simulation using *Rosbridge*[17], which is a JSON-based API providing interoperability with ROS for programs running outside of a ROS environment. More specifically, the JavaScript library *roslibjs*[18], provided by Rosbridge, is used to subscribe to topics published by the main simulation. This includes topics regarding robot states, descriptions, and mission plans as well as information regarding current interventions and task requests.

The monitoring module consists of two parts – a 2D top-down visualization of the simulation environment and a dashboard – which are both briefly introduced in the following subsections.

2D Visualization

The 2D top-down visualization of FleetOR is based on the floor plan of the model environment (see section 4.3.3). As shown in Figure 4.4, this floor plan displays all relevant architectural aspects, such as the locations and dimensions of operating rooms, corridors, peripheral rooms, and doors. Further elements relevant to robotic fleet operation were added, such as a network of robot paths (see section 4.3.3, subsection *Paths*), as used by the fleet manager (see section 4.4) and the robot simulator (see section 4.3.5). Furthermore, the positions of relevant POIs – such as storages, devices, disposals, tube mail stations and robot bases – are visualized. The description of a POI can be displayed by hovering the mouse over the corresponding position. During simulation, the positions of all robots are displayed on the map to visualize the movements of the fleet.

Dashboard

The dashboard summarizes the state of all relevant simulation entities on a single web page, which is shown in Figure 4.5. The first section is dedicated to task requests, providing lists for newly requested tasks, tasks being currently executed and finished tasks. A second section lists all ORs currently in use and provides additional information regarding the surgeon, the type of the intervention being performed, its OPS and ICD codes, and the progress of the surgical workflow. A third section lists all members of the robotic fleet and provides information regarding robot type, capabilities, driving speed, current battery level, current position, current base (if any), availability and the current task being performed. Lastly, the current mission plan of each fleet member is shown, listing all atomic tasks in the planned order. Additional information regarding task type, related task request, score, parameters and target location is given for each atomic task.

[16]https://hub.docker.com/_/httpd
[17]https://wiki.ros.org/rosbridge_suite
[18]https://wiki.ros.org/roslibjs

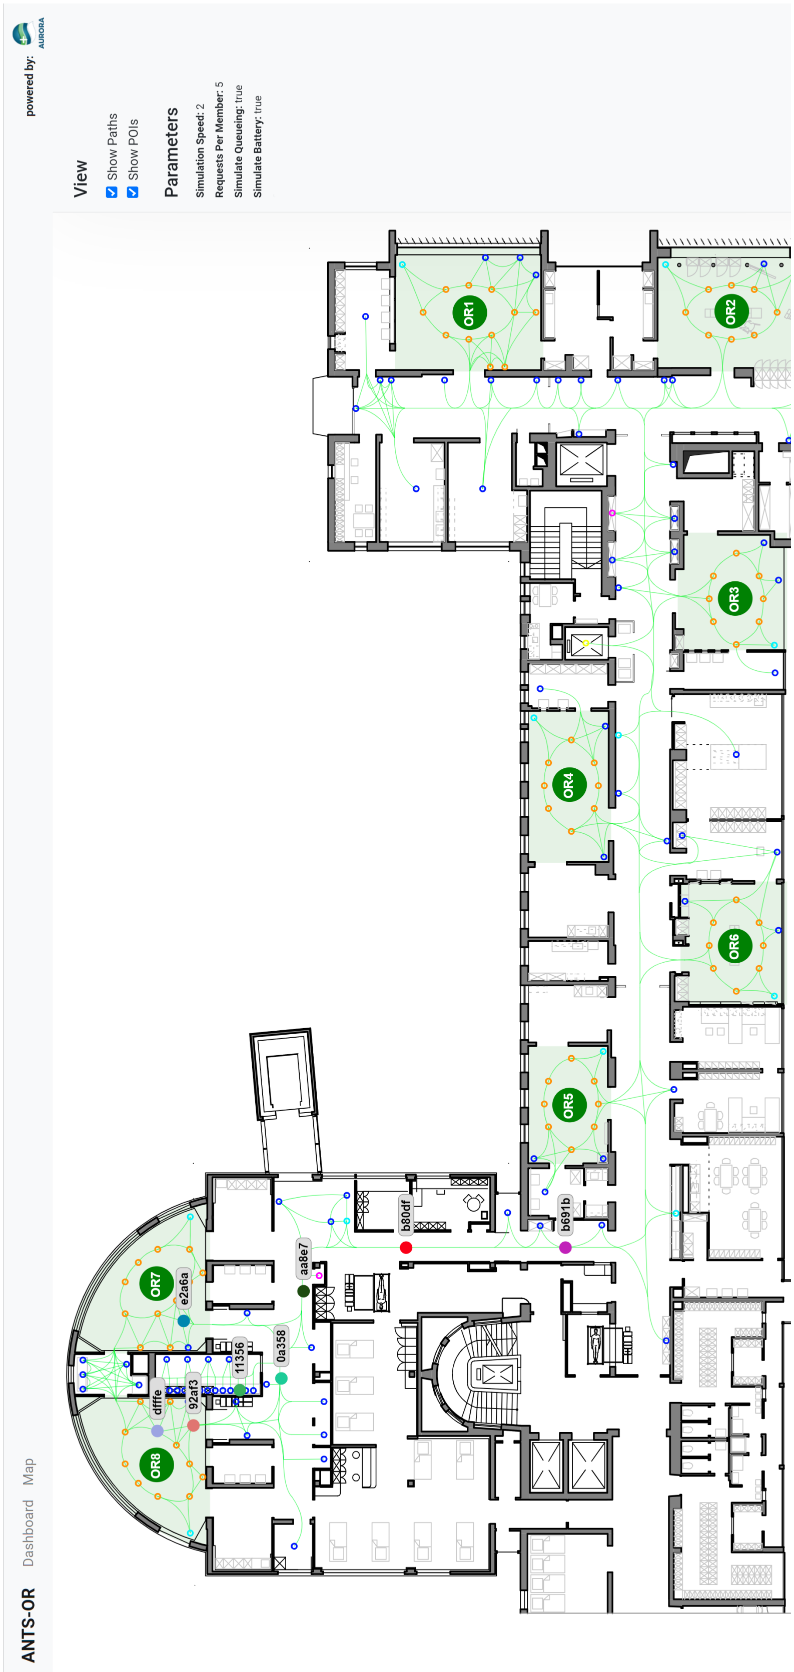

Figure 4.4: The 2D top-down visualization provided by the monitoring module is shown. Within the floor plan of the model OR wing environment, the current positions of the robotic fleet members as well as the locations of POIs are visualized.

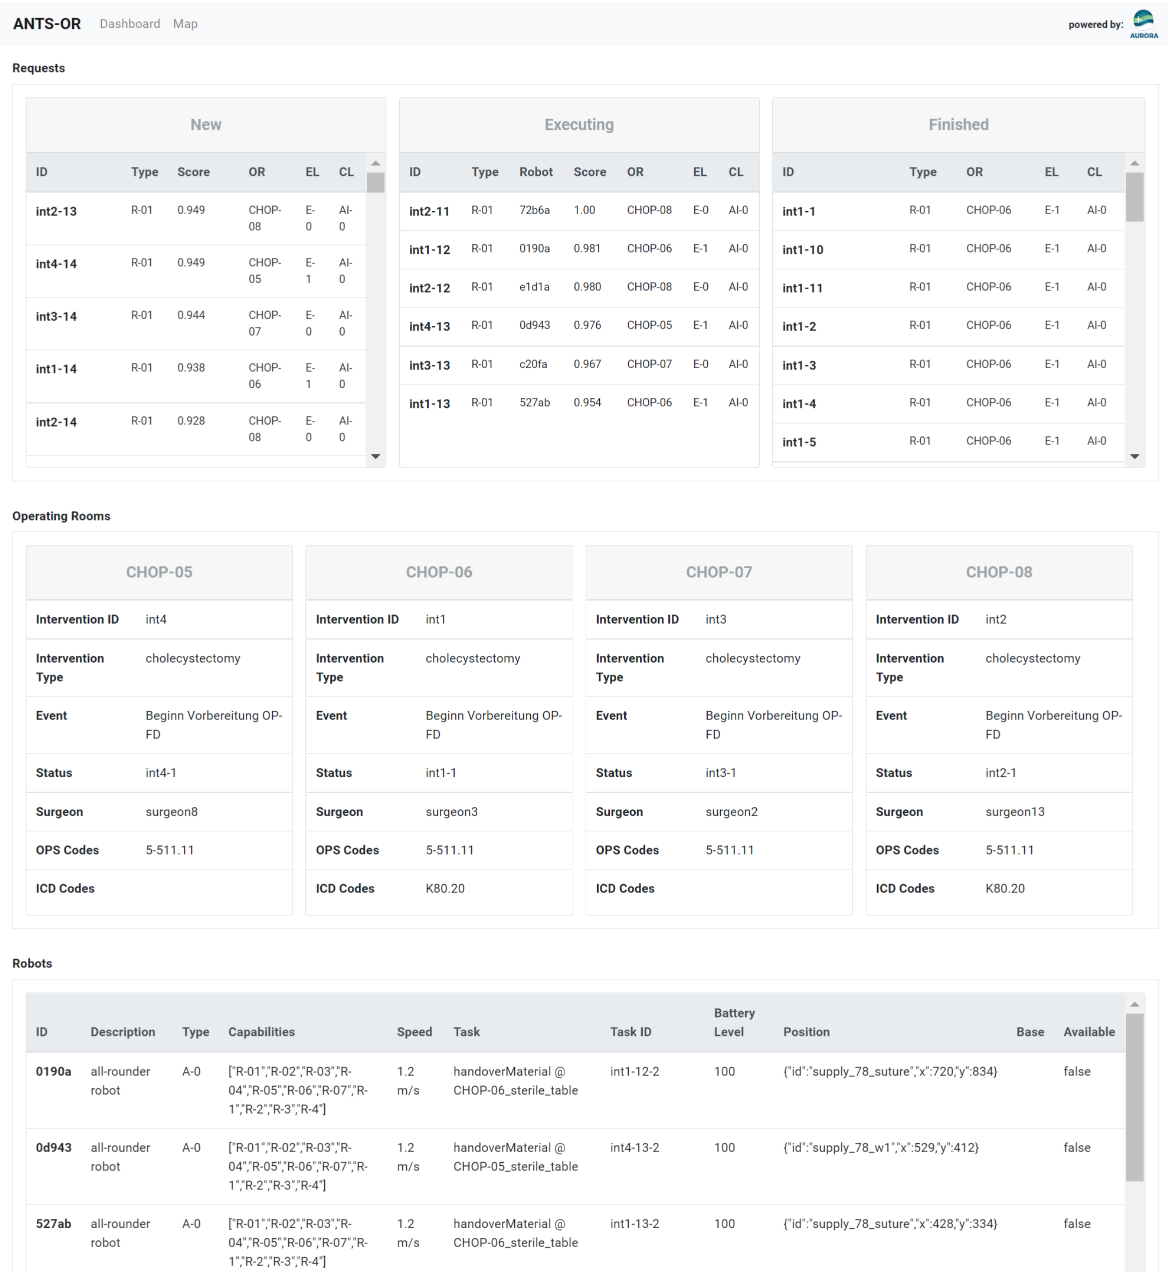

Figure 4.5: The dashboard provided by the monitoring module is shown. Therein, the current state of the simulation can be observed, which includes information regarding currently running surgical interventions, currently executed tasks and robot states.

OR: operating room
ID: unique identifier for tasks, ORs, robots
EL: Emergency Level
CL: Command Level

4.4 Fleet Management

As described in section 2.2, MSR fleet management is concerned with the orchestration of multiple robots, including the allocation of tasks based on incoming requests, the planning of routes through the environment, the integration of new fleet members or the dismissal of faulty or currently abundant robots. Depending on the application scenario, aims and challenges of fleet management may differ considerably.

In the context of this doctoral thesis, MSR fleet management was studied for applications within the operating room wing, which is associated with its own unique characteristics, requirements and quality measures. These aspects are described in section 4.4.1. Subsequent sections cover different sub-problems of fleet management, such as task prioritization (section 4.4.2), task decomposition, and mission planning (section 4.4.3).

4.4.1 Requirements

The integration of MSRs into the OR wing is highly challenging due to hygiene regulations, lack of space, dynamic workflows and safety demands. Since these challenges impact the design of individual robots as well as the concepts used for orchestrating fleets of such systems, they motivate a dedicated analysis, which has been lacking so far. In the following sections, the unique characteristics of the OR wing are described and implications for fleet management are derived.

Hygiene

Hygiene standards are high throughout the OR wing, especially in close proximity to patients currently being treated. Typically, the area surrounding the OR table – including the table itself, the patient and instrument trays – are draped with sterile covers. Persons working within this zone (surgeon, assistant, scrub nurse) are required to wear sterile clothes (scrubs) and must preoperatively perform a defined washing and disinfection routine for hands and forearms. This effort is being made to avoid contamination of the surgical site, which can lead to postoperative complications, such as wound infection or sepsis. With respect to mobile robotics, it can be derived that movements within close proximity to such sterile zones should be carefully slow, or better still, completely avoided. A further hygiene-related concern is the opening of OR doors, which may facilitate the introduction of germs from adjacent rooms. Thus, the number of OR door openings necessary for executing a given collection of task assignments should be minimized by the fleet manager.

Safety

Besides hygiene, other safety aspects are of concern, such as the avoidance of collisions with persons and objects. In general, the OR wing environment is highly challenging in this regard, since available space is scarce and a lot of equipment needs to be placed, interconnected and moved. This is illustrated by Figure 4.6, which shows a typical OR setup for orthopedic surgery. However, space is not only scarce within the operating room but also outside in the corridors and storage rooms, where objects such as medical devices or trays are temporarily stowed.

Figure 4.6: A typical surgical setup used for orthopedic interventions is shown. Clearly, space within the depicted OR is at a premium, since it is equipped with large objects, such as C-arms, monitors and instrumentation tables. Also, there are six members of the surgical team present within the environment.

Clearly, it is therefore desirable to keep the robotic fleet size as small as possible to avoid an over-crowding of the environment. Thus, the fleet size should be carefully dimensioned and composed in order to adhere to space constraints, while ensuring the imposed workload can still be handled acceptably. While the challenges revolving around collision avoidance are more relevant to the development of individual robotic systems, it may have an impact on fleet management as well, since robots are forced to move slowly and additional time is spent for the avoidance of obstacles. In total, these delays may add up and noticeably influence global performance parameters. However, the fleet manager may counter this effect by planning ahead and anticipating upcoming requests. Lastly, the fleet manager should be able to consider safety regulations regarding maximum robot speeds, which may vary among different parts of the OR wing.

Workflow dynamics

Workflows within the OR wing may be highly dynamic and prone to frequent changes or even cancellations. During the perioperative treatment workflow, which involves the collaboration of various clinical entities – such as patient ward, patient transportation service, anesthesia, OR personnel, surgeons, cleaners, technicians etc. – delays may occur for various reasons such as the lack of a signature, the lack of personnel, the lack of materials and so forth. Unexpected events, such as the admittance of urgent emergency cases or unforeseen personnel absences further contribute to this issue. Also, the hard-to-predict nature of surgical interventions themselves poses an imponderability. Beyond that, delays in one workflow may influence the progress of other parallel workflows, since resources (such as operating rooms or wake-up beds) are occupied and cannot be allocated as expected.

Therefore, fleet management concepts for the OR wing must be able to rapidly adapt to changing circumstances. This is particularly relevant for the task allocation process, which assigns incoming task requests to available fleet members. Depending on the problem characteristics, the specific problem instance and the expected solution quality, algorithms for solving such problems are typically computationally taxing, which often results in long processing times. In the OR wing context, however, the fleet manager must be able to re-schedule at arbitrary times and within a very short timespan. According to expert opinions and the findings of the AURORA project, the expected reaction time should be comparable to human reaction times, in particular for emergency situations. This considerably limits the computational complexity of algorithms that are suitable for mission planning. Human reaction times were studied by Deary et al., who report reaction times below 800 ms across all age groups [DD05].

Context Dependence

During surgical workflows, most task requests are expected to be executed as soon as possible after the request has been placed by a surgical team member or by the clinical information system. However, some tasks may be considered more urgent than others, depending on the current context and thus should be executed with preference. Robotic fleet management must be able to consider these differences in urgency, to achieve an adequate and expected fleet behavior depending on the situation at hand.

One central aspect in this regard is the patient condition. In cases where a given task is directly or indirectly linked to the well-being of a patient, it should be executed with priority. As described in [Ber+20], this is especially important for measures dealing with emergency situations, which are common in clinical settings and necessitate immediate response to prevent or mitigate severe outcomes for the patient. Consequently, tasks related to an emergency situation must be prioritized, while non-emergency tasks must be deferred until the critical situation is resolved. Clearly, an objective rating of a task's impact on the health condition of a patient is difficult to impossible, since this involves parameters with subjective or hard-to-predict values. Thus, an ethically fair prioritization is delicate – especially when the well-being of multiple "competing" patients needs to be pondered. This issue is further discussed in section 6.3.

Regarding urgency, the source of a task is usually important contextual information as well. It may originate from a clinical information system, which automatically spawns tasks for the preparation of upcoming surgeries. It may also originate from a person occupying one of many different positions within the clinical command hierarchy. In this context, it is expected that a task requested by an attending physician, for instance, is executed with priority compared to a task requested by a student intern. In general, decisions made by senior team members take precedence over those made by subordinates, as clinical workflows and decision-making processes rely on established hierarchies to facilitate effective collaboration, particularly in urgent and time-sensitive situations [Ber+20]. It is crucial to maintain these command structures in managing MSR fleets by prioritizing tasks requested by senior staff members over those assigned by subordinates [Ber+20].

4.4.2 Task Prioritization

The prioritization of pending orders is a common problem in shopfloor management, where incoming tasks must be allocated to available resources, i.e. machines, in order to optimize global outcome [MLR19]. A similar problem arises in the context of MSR fleet management, where incoming task requests may need to be prioritized in case resources, i.e. robotic fleet members, are scarce. With regard to the OR wing, it is especially desirable to keep the fleet size as small as possible, due to space constraints and hygienic concerns described in the previous. Therefore, by design, situations may frequently arise, where the robotic fleet is not able to execute all pending tasks simultaneously. In such cases, a strategy for determining the relative urgency of tasks is required, which ultimately prescribes the desired order of execution.

However, the outcome of task prioritization heavily depends on the chosen measures of quality. In the clinical context, some of these measures differ considerably from those found in other domains and application scenarios. In industrial manufacturing, for example, economic concerns are usually of the highest relevance. While economic aspects, such as cost efficiency, must be considered in the clinical context as well, other factors, such as patient well-being and patient outcome, are even more important.

The following section describes and motivates two prioritization approaches that were used in the context of the simulation studies presented in the Results chapter.

FIFO

In the most trivial case, tasks can be prioritized with respect to the time of request. This corresponds to a first-in-first-out (FIFO) approach. This strategy is very straightforward to implement and, at the same time, quite natural for the OR wing scenario. Users, i.e. members of the surgical teams, usually expect tasks to be executed as soon as possible after the request has been made. The FIFO approach reflects this "ad-hoc" or "on-demand" workflow style in the sense that the idle time, i.e. the timespan in which the task is waiting for execution, is minimized.

ANTS-OR

However, idle time is not the only relevant quality measure in the context of OR wing workflows – and this is where the FIFO strategy falls short. As described in section 4.4.1, the importance of a given task depends on other factors, most notably the patient condition and the rank of the surgical team member making the request. Including these aspects is expected to yield prioritization results that are closer to what clinical experts would expect – for example, a senior surgeon dealing with a critical bleeding clearly should clearly be served with priority compared to an apprentice nurse preparing an OR for the next upcoming intervention.

Such aspects are considered in the prioritization strategy proposed by the robotic task allocation approach *Auto-navigation Task Scheduling for the OR* (ANTS-OR), which was first described in [Ber+20]. Therein, task requests are prioritized using a three-part weighted score, which can be calculated from *emergency level* L_E, *command level* L_C and *idle time function* f_I:

$$S_{ANTS-OR} = w_E * L_E + w_C * L_C + w_I * f_I(t - t_{request})$$

$S_{ANTS-OR}$:	ANTS-OR Score
L_E:	Emergency Level
w_E:	Weight of Emergency Level
L_C:	Command Level
w_C:	Weight of Command Level
f_I:	Idle Time Function
w_I:	Weight of Idle Time
$t_{request}$:	Time of Request

Herein, the emergency level L_E expresses an urgency from the perspective of the patient. As part of the ANTS-OR scoring concept, the levels shown in Figure 4.7 have been defined. If a given task is requested in the context of a situation where a patient's life is threatened, it should clearly be executed with the highest priority, which is why the maximum emergency level (level 2) is assigned. In situations where an immediate execution of the task is required to avoid damage to the patient's health without fatal consequences, an elevated emergency level (level 1) is assigned. Routine tasks that are not directly connected to the health of patients are assigned with the lowest emergency level (level 0). The influence of the emergency level term within the total $S_{ANTS-OR}$ score can be adjusted using the weight w_E.

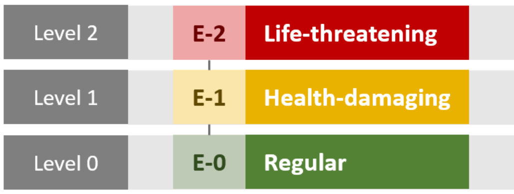

Figure 4.7: The ANTS-OR approach proposes three levels for expressing the urgency of a given task from the perspective of the patient. Non-critical tasks associated with routine situations are assigned the lowest level E-0. In contrast, tasks that are related to the patient's health or may even lead to life-threatening situations if not carried out immediately, are assigned one of the higher priority levels E-1 or E-2. (Figure from [Ber+20])

The command level L_C reflects the position of the requesting person within the clinical command hierarchy. As shown in Figure 4.8, the ANTS-OR concept defines seven distinct levels. The highest level is designated for tasks requested by attending physicians (German: *Oberarzt*) (level 6), who are highly experienced and hold the highest authority in the context of the OR wing. This is followed by surgical fellows (German: *Facharzt*) (level 5), chief nurses (level 4), residents (German: *Assistenzarzt*) (level 3), regular nurses (level 2), digital / AI-based systems (level 1), and medical students or nursing apprentices (level 0). Again, the influence of the command level term within the total $S_{ANTS-OR}$ score can be adjusted using the weight w_C.

In the future, some tasks for MSRs may be requested by clinical information systems or workflow assistance algorithms. For this purpose, a dedicated command level was added (level 1). Such systems could be used for generating preparation tasks, e.g. during the preoperative preparation phase of upcoming surgical interventions. Other sources could be workflow AIs for tracking the progress of an ongoing intervention and automatically spawning assistive task requests to be executed by the robotic fleet. While the requests of clinicians directly express current needs that should be addressed immediately, those made by digital systems are more of a forward-looking nature, which is why their command level is low.

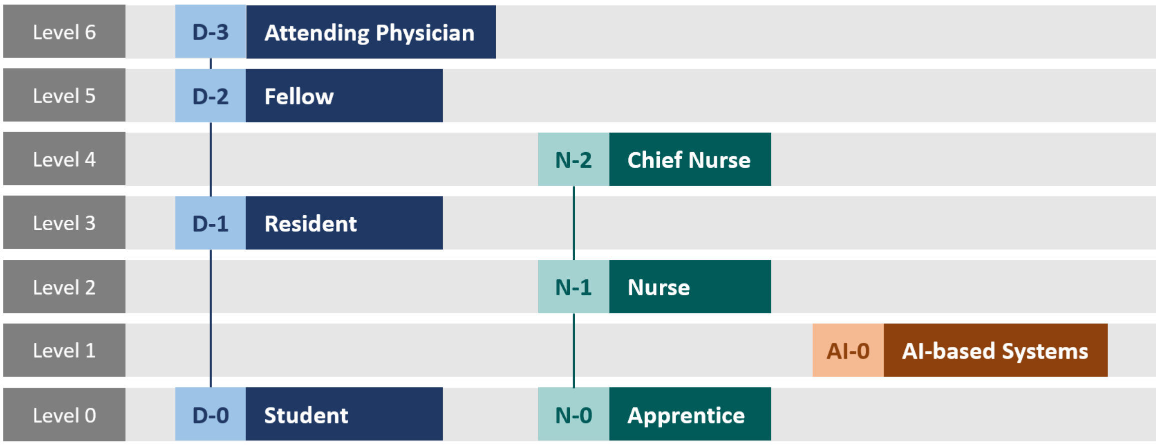

Figure 4.8: ANTS-OR defines seven command levels, which reflect the typical command hierarchy within the hospital. A dedicated level for AI-based systems was added to address the increasing relevance of such systems for the planning and execution of surgical interventions. (Figure from [Ber+20])

The last term of the $S_{ANTS-OR}$ score is the idle time function. Its purpose is to increase the priority of a given task request over time, even if its emergency and command levels are low. This mechanism was added to avoid that the execution of low-priority tasks is completely suppressed in situations with a large number of high-priority tasks. The rationale for this is that, in general, even low-priority tasks can become a hindrance for clinical workflows at some point, if their execution is constantly delayed. In its most basic form, the idle time function is linear in time t:

$$f_I = t - t_{request}$$

The idle time function may also take non-linear forms to, e.g., quadratically or exponentially increase task priority over time:

$$f_I = (t - t_{request})^2$$
$$f_I = e^{t-t_{request}} - 1$$

Such configurations increase the score gently in the short term and more aggressively in the long term. Again, the influence of the idle time function within the total $S_{ANTS-OR}$ score can be adjusted using the weighting factor w_I. For the special case of $w_E = 0$, $w_C = 0$ and a linear idle function f_I, the $S_{ANTS-OR}$ score is reduced to basic FIFO scoring. Thus, $S_{ANTS-OR}$ can be considered an extension of FIFO.

The following exemplary scenario described in [Ber+20] illustrates the ANTS-OR scoring concept:

"*Let's suppose four operating rooms are being in use within an OR wing. The human staff members are supported by a single multifunctional AMR* [Autonomous Mobile Robot] *that is able to fetch supplies as well as operate medical devices. In OR 1, a laparoscopic cholecystectomy (gallbladder removal) is conducted, while the surgical workflow is tracked and assisted by an AI algorithm. Just now, this algorithm has assigned a new task T_1 with the goal of modifying the position of the OR table to prepare for wound closure. In OR 2, the surgical team is currently facing a severe bleeding during a partial hepatectomy. Thus, task T_2 is assigned by the attending* *surgeon with the goal of fetching new blood bags. OR 3 is currently being prepared for an upcoming surgery by the nursing team. The AMR is ordered to reset the surgical devices (task T_3). In OR 4, a port implantation is being performed by a surgical resident under the supervision of an attending surgeon. The resident orders the AMR to fetch suture material (task T_4)."*

The resulting $S_{ANTS-OR}$ scores and the execution order (rank) of the tasks are given in Table 4.4, based on the assumption that all tasks were assigned at the same time. Idle times were disregarded since they are identical for any future point in time due to the simultaneous assignment. The weights w_C and w_E were selected such that both command level and emergency level have an equal impact on the score ($w_C = 1$; $w_E = 3$). As a result, task T_2 receives the highest score, which is reasonable considering it originates from a high-level team member and involves an emergency situation. All other tasks will have to be delayed until T_2 is finished.

Table 4.4: Score values and resulting ranks for tasks T1 to T4 according to the exemplary scenario introduced in [Ber+20].

	L_C	L_E	$S_{ANTS-OR}$	rank
T_1	1	0	1	4
T_2	6	2	12	1
T_3	2	0	2	3
T_4	3	0	3	2

4.4.3 Mission Planning

As described in section 4.3.4, *task requests* represent a complete, clearly defined and delimited assignment to the robotic fleet. While this representation is meaningful from the perspective of the user, it comes with limitations regarding fleet management. Most importantly, treating task requests as encapsulated entities does not allow for an intelligent merging of pending tasks into longer multi-task missions, during which a certain number of tasks is executed along an optimized route through the OR wing. For that, it is beneficial to break down task requests into smaller entities, which are hereinafter referred to as *atomic tasks*. In the context of the mission planning concepts presented in the following, such atomic tasks consist of a *driving* part followed by an *action* part. The description of an atomic task thus defines a target position – to which the robot must navigate first – and an activity to be carried out after this position has been reached. Based on this principle, the task requests defined in section 4.3.4 can be decomposed into atomic tasks according to Table 4.5.

After decomposition, the resulting set of atomic tasks can be reassembled to form overarching missions. The goal is to allocate the decomposed atomic tasks to the available robots, while simultaneously finding driving routes through the OR wing along which the allocated tasks are to be executed. For each robot, a dedicated mission plan is assembled, which may combine atomic tasks originating from multiple requests, as long as the atomic tasks of each single request are executed in the correct order and are all assigned to the same robot. By means of adequate mission planning and routing, a timely execution of high-priority tasks can

Table 4.5: The task types defined in section 4.3.4 can be decomposed into consecutive atomic tasks as follows. Each atomic task requires the robot to drive to a target location and execute a certain action.

Task Type	Atomic Task	Driving	Action
T1	#1	Move to storage	Take out sterile material
	#2	Move to sterile table	Hand over material
	#3	Move to waste disposal	Dispose of packaging
T2	#1	Move to target object	Load / couple object
	#2	Move to target location	Unload / decouple object
T3	#1	Move to target OR	Load waste
	#2	Move to waste disposal	Dispose of waste
T4	#1	Move to target OR	Load sample
	#2	Move to tube mail station	Send sample
T5	#1	Move to target device	Adjust device settings
T6	#1	Move to target OR	Assist sterile clothing

be achieved and driving distances can be reduced, as illustrated by the following example: Consider two task requests of type T1 originating from the same operating room and requiring material from the same storage. Without multi-task mission planning, these tasks would be assigned and executed consecutively, as shown in Table 4.6. However, it is clearly possible to eliminate some of the trips by rearranging the order of execution. This leads to the shortened mission plan shown in Table 4.7, which requires shorter driving distances per request and is also beneficial from a hygienic perspective, since the number of OR door openings is effectively reduced (see section 4.4.1).

Table 4.6: A trivial, i.e. consecutive, mission plan for two tasks of type T1 (task A and task B) is shown. Each task was decomposed into atomic tasks according to Table 4.5. As can be easily observed, the executing robot must move back and forth twice between storage, hand-over location and disposal in order to complete the mission.

Atomic Task	Driving	Action
#A1	Go to storage	Take out sterile material A
#A2	Go to sterile table	Hand over material A
#A3	Go to disposal	Dispose of packaging A
#B1	Go to storage	Take out sterile material B
#B2	Go to sterile table	Hand over material B
#B3	Go to disposal	Dispose of packaging B

Table 4.7: An optimized mission plan for the combined execution of two tasks of type T1 is shown. Here, the atomic tasks are resorted in order to avoid unnecessary trips between storage, hand-over position and disposal.

Atomic Task	Driving	Action
#A1	Go to storage	Take out sterile material A
#B1		Take out sterile material B
#A2	Go to sterile table	Hand over material A
#B2		Hand over material B
#A3	Go to disposal	Dispose of packaging A
#B3		Dispose of packaging B

While this academic example illustrates the benefits of merging multiple requests, real-world mission planning is not as straightforward. The complexity of assembling multi-task missions increases considerably for larger problem sizes, i.e. for a large number of tasks (e.g., 50), a large number of locations within the environment (e.g., 100) and more than one robotic system. Since not all materials can be collected from the same storage, a route may be required to pass through several storage locations before visiting one or more delivery locations. This is illustrated by Figure 4.9.

The described task of planning robot missions along routes through the OR wing can be interpreted as an optimization problem. In the following, the properties of this problem are fully described based on the OR wing application scenario and corresponding requirements. It is further established that this problem is unique to the OR wing use case and has not been studied in scholarly literature before. As one central contribution of this doctoral thesis, a formal model is given and a computationally efficient solution strategy is presented and evaluated.

Problem Formulation

Multi-task mission planning for the OR wing can be understood as an extension of the well-known *Vehicle Routing Problem* (VRP), where least-cost routes must be found for delivering goods from a central depot to customers using multiple vehicles [WN14]. This problem is commonly modeled as an *integer linear program* (ILP), which refers to a class of optimization problems where unknown variables may only take integer solutions and are subject to constraints specified as equalities and inequalities [CCZ14]. Solving ILPs is NP-hard, meaning that the required computational effort increases exponentially when increasing the size of the problem instance (e.g., more vehicles, more customers) [CCZ14]. As of yet, no methods exist for finding optimal solutions to such problems within polynomial time (i.e., with polynomially, not exponentially increasing computational effort) [KV07]. Therefore, the available optimal solving strategies are of limited value for practical applications with larger-scale problem instances. Instead, sub-optimal strategies are commonly used and have provided very good results for many applications in the past.

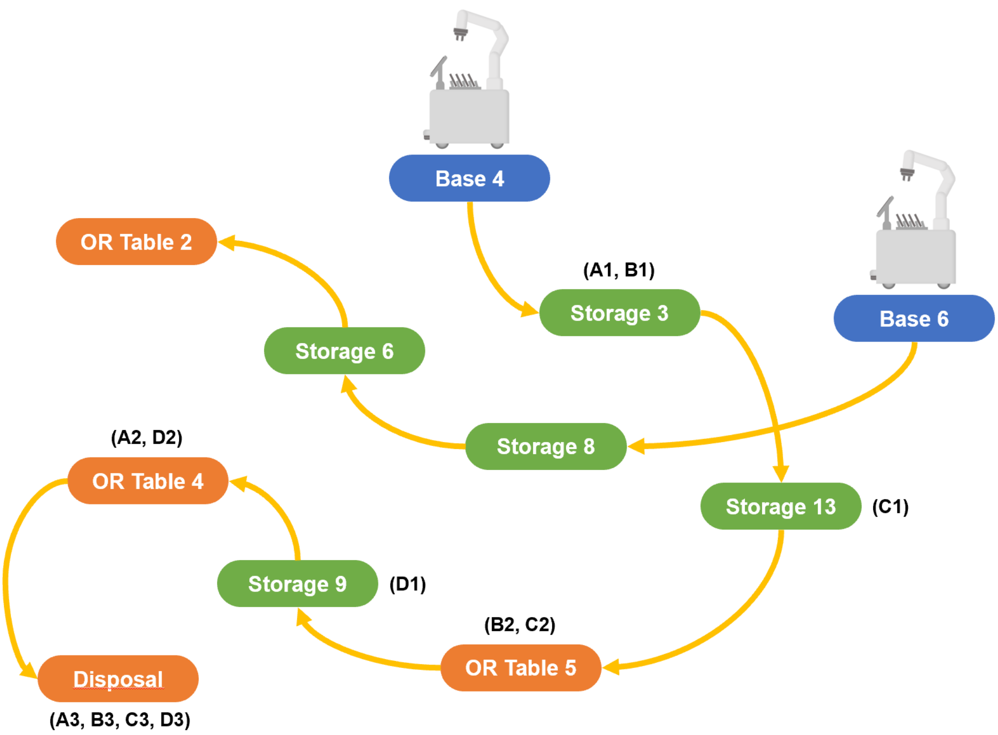

Figure 4.9: Exemplary mission routes are shown for two robots. Starting from one of the robot bases, each robot visits different locations within the environment while executing atomic tasks according to its mission plan, as indicated in parentheses for one of the routes. By planning routes such that multiple atomic tasks are executed at the same location wherever possible, driving distances and thus task execution durations can be reduced.

Based on the formulation in [Bor17], the classical VRP can be modelled as an ILP using the following parameters:

- **N**: Set of vehicles
- **L**: Set of locations (customers and depot)
- $\mathbf{s_i}$: Demand of customer i
- $\mathbf{s_{max}}$: Capacity of each vehicle
- $\mathbf{c_{i,j}}$: Cost for travelling from location i to location j

The routes, which represent the solution of the optimization problem, are described using a three-indexed binary decision variable:

$$\mathbf{x_{n,i,j}} = \begin{cases} 1 & \text{if vehicle } n \text{ moves from location } i \text{ to location } j \\ 0 & \text{else} \end{cases} \tag{4.1}$$

The optimization goal is to find a route for each vehicle, such that overall driving costs are minimized, which can be modelled as follows:

$$min \sum_{n \in N} \sum_{i \in L} \sum_{j \in L} c_{i,j} x_{n,i,j} \tag{4.2}$$

Furthermore, the following constraints must be satisfied by the solution:

- Each customer is visited by exactly one vehicle:

$$\sum_{n \in N} \sum_{i \in L} x_{n,i,j} = 1 \qquad \forall j \in L, \ i \neq j \tag{4.3}$$

- The central depot may only be left once by each vehicle:

$$\sum_{j \in L} x_{n,0,j} = 1 \qquad \forall n \in N \tag{4.4}$$

- The number of arrivals and departures at each location must be equal:

$$\sum_{i \in L} x_{n,i,j} = \sum_{j \in L} x_{n,j,i} \qquad \forall n \in N, \ \forall j \in L, \ i \neq j \tag{4.5}$$

- The sum of demands of all customers visited by a vehicle must be smaller than or equal to the vehicle's capacity:

$$\sum_{i \in L} \sum_{j \in L} s_j x_{n,i,j} \leq s_{max} \qquad \forall n \in N, \ i \neq j \tag{4.6}$$

- Subroutes disconnected from the depot are not allowed (subtour elimination constraints):

$$\sum_{i \in S} \sum_{j \in S} x_{n,i,j} \leq |S| - 1 \qquad \forall S \subseteq \{2, ..., |L|\}, \ \forall n \in N \tag{4.7}$$

While the classical VRP introduced above provides a baseline and a general paradigm for problem formulation, it is far from sufficient for addressing the requirements of multi-task robotic mission planning within the OR wing. For adapting the VRP accordingly, several fundamental extensions must be made, necessitated by the characteristics of the application scenario:

- **Multiple Depots:** Robot missions may start at arbitrary locations (depots), i.e. either from one of the home bases or from the end location of the last mission. This is referred to as the *Multi-Depot Vehicle Routing Problem* (MDVRP) [RGP20].
- **Open Routes:** Missions are not required to end at the same location (depot) where they were started, which is referred to as the *Open Vehicle Routing Problem* (OVRP) [SW19].
- **Pickups and Deliveries:** As described in section 4.1.1, some task types may require the pickup of an object at one location and the delivery to another location (e.g., the delivery of a sterile material from a storage to the OR table). This is referred to as the *Vehicle Routing Problem with Pickups and Deliveries* (VRPPD), in the variant described in [Ber+07]. Since a material must first be picked up before it can be delivered, precedence constraints are introduced.
- **Alternative Pickups:** Due to the redundant storage concept commonly found within the OR wing, some materials may be available at multiple storages. Therefore, alternative pickup locations must be considered when planning robot routes. The aim is to select the optimal alternative with respect to given quality measures (usually driving duration). To the best of the author's knowledge, this aspect has not yet been considered by other work. One related problem referred to as the *Capacitated Vehicle Routing Problem with Alternative Delivery, Pickup and Time Windows* was studied by Sitek et al. [Sit+21]. However, the work focuses on alternative delivery locations, while not considering alternative pickup locations.
- **Task Priorities:** In contrast to the classic VRP, tasks may have different priorities, which must be considered when assembling missions. Tasks of high priority should be executed with preference, while it is acceptable to delay the execution of low-priority tasks.
- **Route Length**: If robot missions become too long, the reaction time of the system is poor when dealing with spontaneously requested tasks of high priority. Instead, it is reasonable to define an upper limit for the mission length or to flexibly break down long missions into multiple shorter ones.
- **Heterogeneous Fleet:** In the general case, the robotic fleet may consist of multiple robot types with varying capabilities. A given robot type may either be able to execute all or only some types of tasks.

The goal of robotic fleet management within the OR wing is the composition of robot missions from pending task requests, while fulfilling the above requirements and optimizing for driving duration and task priority. In the following, this novel VRP variant will be referred to as the *Vehicle Routing Problem for the OR Wing* (VRP-OR). With respect to the taxonomies summarized in section 2.2, the VRP-OR can be classified as *MT-SR-TA:SP* (multi-task robots, single-robot tasks, time-extended assignment with precedence constraints) [GM04; Nun+17], subject to *in-schedule dependencies (ID)* [KSD13].

With reference to the formulation of the classical VRP given in the previous, a mathematical model of the VRP-OR is provided in the form of an integer linear program. Therein, robots may be heterogeneous with respect to the types of tasks they are capable of executing, the number of items they can carry and their driving speed. It is assumed that a given task request cannot be decomposed into atomic tasks in more than one way. For some tasks, alternative execution locations may exist and finding the optimal choice is among the objectives of the optimization problem. The following parameters are used in the model formulation:

- $\mathbf{R}$: Set of pending requests
- $\mathbf{D_r}$: Decomposition of request r, i.e. the set of atomic tasks constituting this request
- $(\mathbf{t} \rightarrow \mathbf{u})$: Precedence constraint specifying that atomic task t must be executed before atomic task u (not necessarily directly). Such constraints are introduced when decomposing a request into atomic tasks.
- $\mathbf{T}$: Set of all atomic tasks decomposed from the pending requests, i.e. $T = \bigcup_{r \in R} D_r$ and thus $D_r \subset T, \ \forall r \in R$
- $\mathbf{h_t}$: Capability required for executing task t
- $\mathbf{p_t}$: Priority of task t with $p_t \in [0; 1]$
- $\mathbf{z}$: Penalty for not assigning a task
- $\mathbf{N}$: Set of robots
- $\mathbf{H_n}$: Set of capabilities of robot n
- $\mathbf{v_n}$: Driving speed of robot n
- $\mathbf{s_{max,n}}$: Inventory size of robot n, denoting the number of slots available within the inventory of robot n
- $\mathbf{s_t}$: Number of inventory slots required for executing task t
- $\mathbf{m_{max}}$: maximum mission length, i.e. maximum number of tasks that may be assigned to a robot.
- $\mathbf{B}$: Set of base locations
- $\mathbf{b_n}$: Base of robot n at which it starts its mission, with $b_n \in B$
- $\mathbf{A}$: Set of task locations
- $\mathbf{L}$: Set of all locations, i.e. $L = A \cup B$
- $\mathbf{d_{i,j}}$: Distance between locations i and j, $\forall i, j \in L$[19]
- $\mathbf{A_t}$: Set of alternative execution locations of task t, with $A_t \subset A$

Two three-indexed binary decision variables are introduced for describing the path of each robot and the tasks executed along the way:

[19] By defining a distance of 0 from any task location to any robot base (but not in the other direction), the problem becomes an open-loop VRP. Otherwise, robot routes are closed loops, as typically required in the context of the classical VRP.

$$\mathbf{x_{n,i,j}} = \begin{cases} 1 & \text{if robot } n \text{ moves from location } i \text{ to location } j \\ 0 & \text{else} \end{cases} \tag{4.8}$$

$$\mathbf{y_{n,t,i}} = \begin{cases} 1 & \text{if atomic task } t \text{ is executed by robot } n \text{ at location } i \\ 0 & \text{else} \end{cases} \tag{4.9}$$

Furthermore, the optimization problem is subject to the following constraints, which are based on the classical VRP and on further requirements introduced by the OR wing application scenario:

- Each task may be executed at most once:

$$\sum_{n \in N} \sum_{i \in L} y_{n,t,i} \leq 1 \qquad \forall t \in T \tag{4.10}$$

- The sum of arrivals and departures to/from task locations must be equal:

$$\sum_{i \in L} x_{n,i,j} = \sum_{i \in L} x_{n,j,i} \qquad \forall n \in N,\ \forall j \in T,\ i \neq j \tag{4.11}$$

- Each task location must be visited by the robot:

$$y_{n,t,j} \leq \sum_{i \in L} x_{n,i,j} \qquad \forall n \in N,\ \forall t \in T,\ \forall j \in L \tag{4.12}$$

- If no task is executed at a possible task location, it is not visited by the robot:

$$\sum_{i \in A} x_{n,i,j} \leq \sum_{t \in T} y_{n,t,j} \qquad \forall n \in N,\ \forall j \in A \tag{4.13}$$

- All (atomic) tasks belonging to the decomposition D_r of a given request must be assigned to the same robot:

$$\sum_{t \in D_r} \sum_{i \in A} y_{n,t,i} = |D_r| \sum_{i \in A} y_{n,u,i} \qquad \forall n \in N,\ \forall r \in R,\ \forall u \in D_r \tag{4.14}$$

- Each robot must leave its base exactly once:

$$\sum_{j \in A} x_{n,i,j} = 1 \qquad \forall n \in N,\ i = b_n \tag{4.15}$$

- Bases may not be visited during the mission:

$$\sum_{b \in B} \sum_{j \in L} x_{n,b,j} = 0 \qquad \forall n \in N,\ b \neq b_n \tag{4.16}$$

- The number of tasks assigned to a robot is limited by the maximum mission length:

$$\sum_{t \in T} \sum_{i \in L} y_{n,t,i} \leq m_{max} \qquad \forall n \in N \tag{4.17}$$

- The number of items a robot can store is limited by the slots of its inventory:

$$\sum_{t\in T}\sum_{i\in L} s_t y_{n,t,i} \leq s_{max,n} \qquad \forall n \in N \tag{4.18}$$

- A robot must be able to execute an atomic task for it to be assignable:

$$\sum_{i\in L} y_{n,t,i} = 0 \qquad \text{if} \qquad h_t \notin H_n \qquad \forall n \in N,\ \forall t \in T \tag{4.19}$$

- Each task may only be executed at one of its given alternative locations:

$$y_{n,t,i} = 0 \qquad \text{if} \qquad i \notin L_t \qquad \forall n \in N,\ \forall t \in T,\ \forall i \in L \tag{4.20}$$

- Precedence constraints[20]:

$$\sum_{k\in A} x_{n,k,j} - \sum_{k\in A} x_{n,k,i} \geq 0 \qquad \forall n \in N,\ \forall (t \rightarrow u),\ \forall i \in A_t,\ \forall j \in A_u \tag{4.21}$$

- Subtour elimination constraints:

$$\sum_{i\in S}\sum_{j\in S} x_{n,i,j} \leq |S| - 1 \qquad \forall S \subseteq \{2, ..., |A|\},\ \forall n \in N \tag{4.22}$$

The quality function aims at minimizing driving durations and enforces a priority-dependent penalty for tasks that are not assigned:

$$min \sum_{n\in N}\sum_{i\in L}\sum_{j\in L} \frac{d_{i,j}}{v_n} x_{n,i,j} + p_t z(|T| - \sum_{n\in N}\sum_{t\in T}\sum_{i\in L} y_{n,t,i}) \tag{4.23}$$

Solution Strategy

Since all VRP variants are extensions of the *Traveling Salesperson Problem*, finding the optimal solution is NP-hard [CK21; Sit+21], meaning that the required computation time increases exponentially when increasing the size of the problem instance (e.g. by increasing the fleet size or the number of task requests). Consequently, this is also true for the VRP-OR. At the same time, the characteristics of OR wing workflows require fast rescheduling times, ideally comparable to human reaction times, as motivated in section 4.4.1. Thus, optimal solution methods, such as exhaustive search or the branch-and-bound method [Lit+63], are clearly not usefully applicable, except for very small problem instances. Instead, heuristic methods are commonly used when instances of NP-hard problems must be solved within shorter time-spans. Heuristics do not necessarily provide the optimal solution, but often lead to a "good-enough" result that is acceptably close to the optimum.

To that end, a greedy task allocation algorithm is proposed for heuristically solving the VRP-OR. Greedy algorithms aim at finding a locally optimal solution by repeatedly making the locally optimal decision, which may or may not lead to the global optimum [Vin02]. As

[20] It is assumed that there is no overlap in the locations for pickups (set A_p) and the locations for deliveries (set A_d), i.e. $A_p \cap A_d = \emptyset$ and $A_p \cup A_d = A$. It is further assumed that only the preceding task (t) has alternative locations. Both assumptions hold in the context of the use cases studied within this doctoral thesis.

one major benefit, greedy algorithms exhibit a low computational complexity and thus allow for short processing times. This is in line with the requirements of the OR wing environment, where the behavior of the fleet must be quickly adapted to changing circumstances.

According to the proposed algorithm, task requests are assigned sequentially, based on their priority rating. Following the greedy paradigm, the mission plan resulting in the lowest (local) cost for the currently processed task is selected at each iteration, while neglecting the influence of the remaining lower-priority tasks that must be assigned in subsequent iterations. While this comes at the price of foregoing strict optimality, a solution can be found in polynomial time, which is required for achieving acceptably short rescheduling durations.

The proposed algorithm consists of the following steps: First, all pending task requests from the surgical team (or the clinical information system) are prioritized according to the selected prioritization strategy (e.g. FIFO or ANTS-OR as described in section 4.4.2).

Before these prioritized tasks can be assigned to robots, they must first be decomposed into atomic tasks:

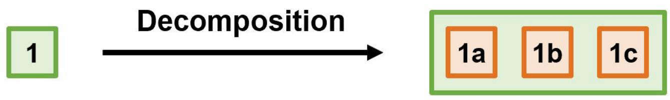

Starting from the top of the prioritized task list, the mission planner now assigns tasks to members of the robotic fleet. In each case, it is first checked whether the robot is capable of executing this type of task and whether the remaining load capacity of the robot is sufficient. For the purpose of task assignment, each robot possesses a mission plan of configurable length, which may already contain some previously assigned atomic tasks.

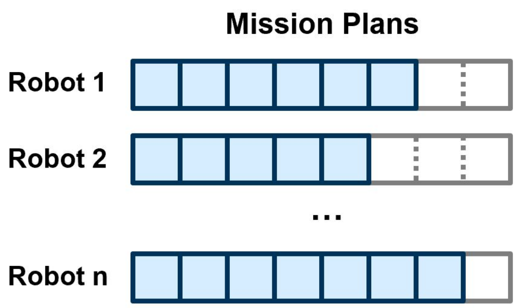

The first atomic task of the decomposed request with highest priority is now tentatively inserted at all possible positions within the current mission plan of each robot. This creates a set of alternative mission plans, with the atomic task inserted at different positions.

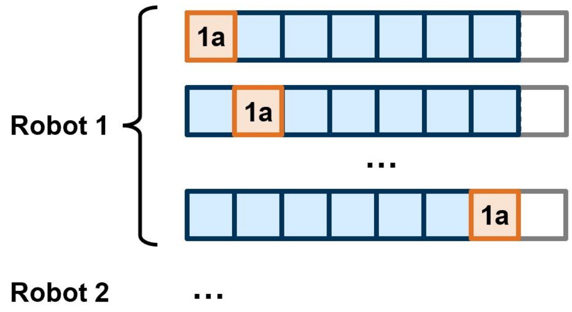

For each of the generated mission alternatives, the same process is now repeated with the second atomic task of the request, and so forth. Since there are precedence constraints between atomic tasks that belong to the same request, the $(k+1)$-th atomic task of the request must be inserted at some position after the k-th atomic task.

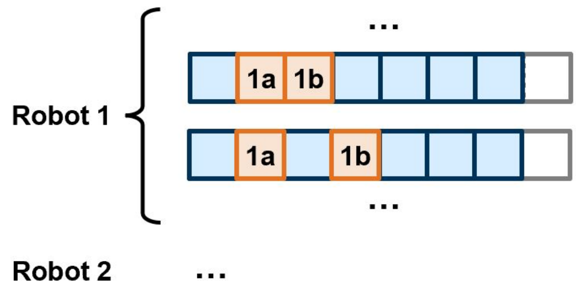

Using the described approach, possible sequences of atomic tasks are generated, which are referred to as *mission candidates*. With the number of atomic tasks within the decomposition of a given task t denoted as $|D_t|$, the set of robots N and the previous mission length $\widetilde{m}_n$ of robot n, the number of mission candidates C is:

$$C = \sum_{n \in N} (\widetilde{m}_n + 1)^{|D_t|} \tag{4.24}$$

Thus, C increases exponentially in $|D_t|$, however, since $|D_t|$ is typically very small (see Table 4.5), the number of mission candidates remains acceptably low. This would change when permutating the previously assigned atomic tasks as well, in order to truly obtain all possible combinations of atomic tasks, which would be required for exhaustively searching for the globally optimal mission plan. In this case, the number of mission candidates would increase factorially, which quickly leads to extreme numbers of mission candidates, even for smaller problem instances. Instead, the algorithm aims at finding the best position of a given task within the *existing* sequence of previously assigned tasks.

For each mission candidate, multiple driving routes may exist that are suitable for executing this exact sequence of atomic tasks. As described in the previous, this is due to the redundant storage concept of the OR wing, where the same type of material may be available from multiple storages. When increasing the mission length m, the number of possible routes R for a given mission candidate increases exponentially, with $|A_k|$ being the number of location options for the k-th atomic task of the route:

$$R_{optimal} = \prod_{k=1}^{m} |A_k| \tag{4.25}$$

For the special case of an equal number a of location options for all atomic tasks (i.e., $a = |A_k|, \ \forall k = 1, ..., m$), this simplifies to $R = a^m$, which directly shows the exponential behavior in m. Thus, exhaustively searching all possible routes for all mission candidates is not reasonable, especially for increasing mission lengths. Again, a greedy strategy is proposed to obtain a suitable (if not necessarily optimal) route with acceptable computational effort: Following the greedy principle, the final route is generated by cycling over all atomic tasks and repeatedly selecting the shortest distance from the current atomic task to the next one.

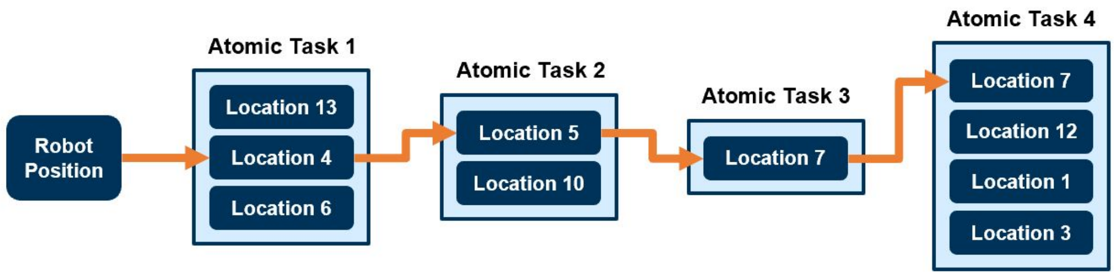

This reduces the number of options that must be searched to:

$$R_{greedy} = \sum_{k=1}^{m} |A_k| \tag{4.26}$$

For an equal number of location options, this simplifies to $R_{greedy} = ma$, which exhibits a linear computational complexity. After finding the greedy route for each mission candidate, the candidate resulting in the shortest driving duration is selected as the final mission plan, which is assigned to the corresponding robot. The entire process described above is repeated until one of two stop criteria is fulfilled: Either the mission plans of all robots are completely filled or there are no further pending task requests that could be assigned. In case no task is assigned to a given robot, it is sent to the nearest home base for recharging.

5

Results

In the following chapter, the conducted simulation-based studies and the obtained results are presented. This pertains to aspects of fleet dimensioning – such as fleet size (section 5.1), fleet composition (section 5.2), robot velocity (section 5.3) and battery life (section 5.4) – as well as aspects regarding fleet management – such as task prioritization (section 5.5), multi-task mission planning (section 5.6), quick storages (section 5.7) and robot inventories (section 5.8). In a final simulation scenario, the findings of these isolated studies are jointly applied and compared to the human-only status quo without robotic support. While the overarching discussion of the results is the subject of the next chapter 6, a short discussion section is already provided alongside the results of each study, for better comprehensibility.

5.1 Fleet Size

One fundamental question in the context of MSR fleets for the OR wing is how the fleet size should be dimensioned, i.e. what number of fleet members is reasonable for a given scenario. As explained in previous sections, a general characteristic of the OR wing environment is that space is at a premium, which is why small fleets are advantageous in this regard. This is in line with economic considerations, for which the avoidance of costly investments, such as MSRs and the associated operational and maintenance costs, is favorable. On the other hand, the fleet must be able to cope with everyday task workload without being constantly overloaded. Therefore, a compromise must be found, which represents the sweet spot for real-world scenarios. For this purpose, a simulation-based study was conducted and the findings were used to derive tangible recommendations for future real-world application. This study was also published in [Ber+22].

Study Design

Clearly, the optimal fleet configuration for a given scenario depends on different parameters, in particular the number of operating rooms that need to be served, the achievable robot velocities and the capabilities of the individual fleet members. Varying these parameters resulted in 144 different scenarios, which were simulated using the FleetOR simulation framework (see section 4.3). The studied scenarios can be categorized into 12 groups, which

Table 5.1: The table defines 12 groups (G1-G12) of simulated scenarios, categorized by robot types, driving speeds, and the number of operating rooms in concurrent usage. Within each group, fleet sizes ranging from 1 to 20 robots (all-rounder robots), or alternatively, 4, 8, 12, 16, and 20 robots (specialized robots), were simulated. (Table from [Ber+22])

	G1	G2	G3	G4	G5	G6	G7	G8	G9	G10	G11	G12
Robot Types	all-rounder robots (R5)						specialized robots (R1-R4)					
Driving Speed v_r	0.3 m/s			1.2 m/s			0.3 m/s			1.2 m/s		
Number of ORs	2	4	6	2	4	6	2	4	6	2	4	6

are summarized in Table 5.1. Each group is associated with a specific combination of robot types, robot driving speeds and number of operating rooms.

The parameter *robot types* refers to the capabilities of individual fleet members: The fleet was either composed of identical all-rounder robots (scenario groups G1-G6) or of specialized robots with differing capabilities (scenario groups G7-G12). Per definition, all-rounder robots are able to execute any type of task that may be requested. In contrast, specialized robots are only able to execute certain tasks. The robot types used for the conducted simulations are defined in Table 4.3.

For scenario groups G1-G3 and G7-G9, the robot driving speed v_r was set to 0.3 m/s for each fleet member and for scenario groups G4-G6 and G10-G12 to 1.2 m/s. The choice of these values was based on the speed thresholds defined by DIN EN ISO 3691-4 (appendix A.2) [DIN20].

In order to eliminate the influence of other factors, single-task scheduling and FIFO prioritization were used for the purposes of this study. While more advanced scheduling methods can help further improve key performance parameters of the robotic fleet, they introduce additional degrees of freedom to the system, which may complicate the interpretation of the results. Instead, the presented study aimed at providing a clean baseline regarding the influences of fleet size, fleet composition and robot speeds. This starting point was later used as a reference for further investigations regarding fleet management, which are presented in subsequent parts of this thesis.

Results

Figure 5.1 shows the total durations of all simulated scenarios. It can be observed that an increasing fleet size leads to a rapid exponential-like decrease of the scenario duration until a saturation level is reached. However, durations may rise again after a turnaround point is passed (e.g. 12 robots for G1), which is the case for scenario groups with two ORs (scenario groups G1, G4, G7, G10). Generally, fleets composed only of all-rounder robots resulted in shorter scenario durations compared to fleets composed of specialized robots. An exception to this are scenario groups G1 and G2, which show a slightly better performance for fleet sizes greater than 12, and scenario groups G4 and G5, which show a slightly better performance for fleet sizes greater than 8. It must be noted that most robotic scenarios were not able to achieve a shorter total duration than the recorded human reference (1-2 dedicated circulating nurses per OR). The sole exceptions to this occur within scenario group G4 (fleet sizes between 3 and 12 robots), and within scenario group G10 (fleet sizes of 8 and 12 robots).

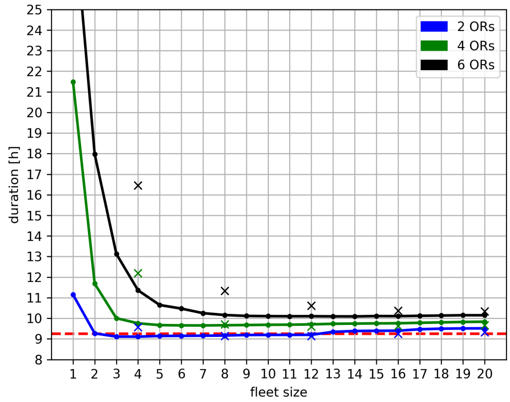

(a) $v_r = 1.2m/s$

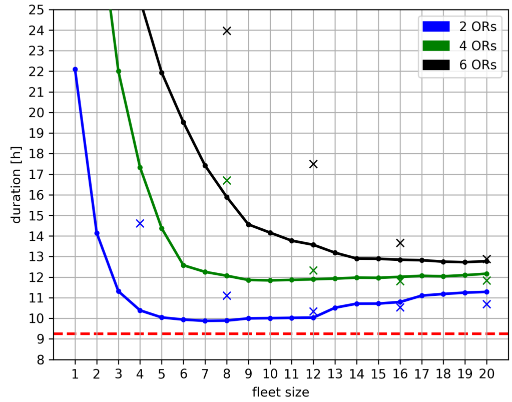

(b) $v_r = 0.3m/s$

Figure 5.1: The plots depict simulation results for total scenario durations (in hours). In part a), robot speed was $v_r = 1.2m/s$, while in part b) it was $v_r = 0.3m/s$. Continous graphs are used for all-rounder robots (groups G1-G6), while individual crosses mark results for specialized robots (G7-G12). The duration achieved for the human-only scenario is depicted as a dashed orange line. (Figures from [Ber+22])

Figure 5.2 shows average task response times. Again, it can be observed that an increase in fleet size leads to a decrease in task response time until a saturation level is reached. Scenarios with specialized robots consistently result in longer response times than scenarios with all-rounder robots. Since intra-operative tasks are not prioritized by the FIFO scheduling used for the simulations, the average of all executed tasks is given. This includes tasks for preoperative preparation and postoperative cleanup.

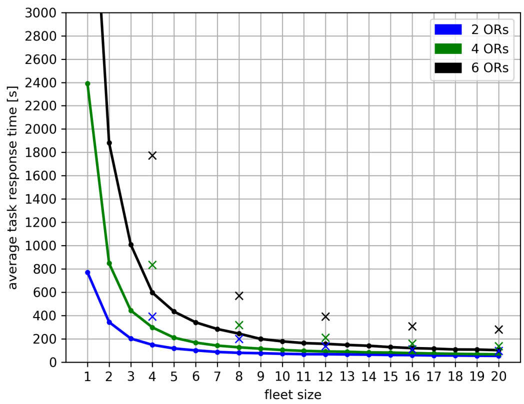

(a) $v_r = 1.2m/s$

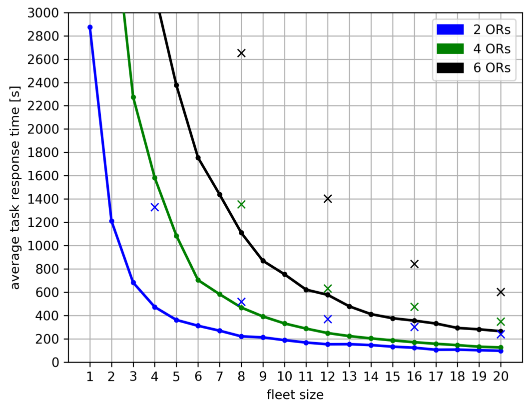

(b) $v_r = 0.3m/s$

Figure 5.2: The plots depict average task response times (in seconds) for the simulated scenarios. (Figures from [Ber+22])

Figure 5.3 illustrates fleet utilization across all simulated scenarios. As anticipated, an expansion in fleet size corresponds to a reduction in fleet utilization, since the total workload remains the same. Notably, an inverse-exponential trend emerges for higher robot speeds and fewer operating rooms, while slower robot speeds combined with a high number of ORs result in an approximately inverse-linear relationship. Furthermore, scenarios featuring specialized robots consistently demonstrate lower fleet efficiencies compared to those with all-rounder robots.

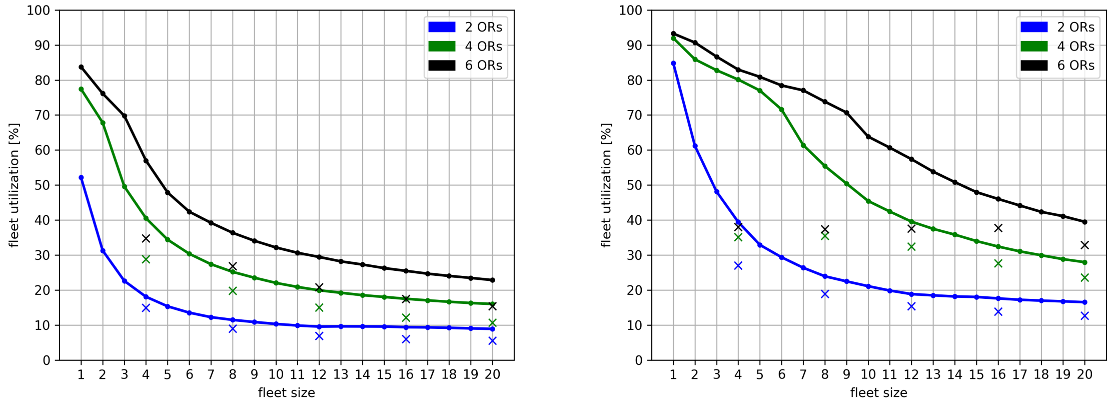

(a) $v_r = 1.2m/s$ **(b)** $v_r = 0.3m/s$

Figure 5.3: The plots depict fleet efficiencies (in percent) for the simulated scenarios. (Figures from [Ber+22])

Discussion

According to the results obtained, it can be established that expanding the size of the fleet does not necessarily lead to improved performance. Although there is a clear limitation to realistically applicable fleet sizes imposed by the available space for driving and parking, other factors come into play even before reaching this limit. Beyond a certain point, increasing the fleet size does not result in further performance enhancement. In fact, for scenarios with two operating rooms (G1, G4, G7, and G10), there is even an inversion of total scenario durations that can be observed.

This phenomenon can be attributed to increased driving times, as the two operating rooms in the associated scenarios are situated on one side of the OR wing (see OR 1 and 2 in Figure 4.3) and bases close to these ORs will be fully occupied quite quickly. As a result, some robots are forced to use more distant bases, causing longer driving times. However, for scenarios where task execution locations are better distributed, the probability of having a robot available at a nearby base is higher.

Based on these findings, the optimal number of robots for an operating room can be determined, i.e., the minimum fleet size that still results in an adequate level of performance. For the purpose of this discussion, a total scenario duration that is at most 5 % longer than the shortest total duration within the same scenario group (same number of ORs, same robot velocity) will still be considered acceptable. For scenarios employing all-rounder robots, this translates to an optimal fleet size of 1 robot per operating room for $v_r = 1.2m/s$ (scenario groups G4-G6) and 2-3 robots per operating room for $v_r = 0.3m/s$ (scenario groups G1-G3). However, the utilization of specialized robots necessitates larger fleet sizes, even for scenarios with fewer operating rooms, resulting in 2 robots per operating room for $v_r = 1.2m/s$ (scenario groups G10-G12) and 3-6 robots per operating room for $v_r = 0.3m/s$ (scenario groups G7-G9). Recommended fleet sizes for each scenario group are given in Table 5.2.

When considering human performance as a benchmark (represented by the orange dashed line in Figure 5.1), only a few robotic scenarios demonstrate superior outcomes. For scenar-

Table 5.2: Based on the simulation results, recommended fleet sizes are given for all scenario groups. (Table from [Ber+22])

	G1	**G2**	**G3**	**G4**	**G5**	**G6**	**G7**	**G8**	**G9**	**G10**	**G11**	**G12**
Robot Types	all-rounder robots (R5)						specialized robots (R1-R4)					
Driving Speed v_r	0.3 m/s			1.2 m/s			0.3 m/s			1.2 m/s		
Number of ORs	2	4	6	2	4	6	2	4	6	2	4	6
Recommended fleet size	5	7	13	2	3	6	12	12	20	4	8	12

ios involving slow robot velocities or more than two operating rooms, the overall duration exceeds that of the human reference across all fleet sizes. This could be attributed to longer travel paths and potential queues. Additionally, it is important to note that the implemented workflow simulation incorporates a waiting period in cases where the robots complete all preoperative preparation tasks for a particular procedure earlier than the recorded human reference. In this regard, the simulations are quite conservative.

When considering other types of surgical interventions, the amount and frequency of tasks may differ depending on various factors, such as the workflow complexity and the degree to which tasks can be prepared preoperatively. Consequently, the optimal number of robots may differ as well. However, the results obtained for the reference intervention (see section 4.1.2) provide a meaningful baseline.

5.2 Fleet Composition

Based on the study presented in section 5.1, conclusions regarding fleet composition can be drawn as well: Throughout the simulation results, it was consistently observed that all-rounder robots outperformed specialized robots. E.g., for scenario group G6 and a fleet size of 4, an approximately 31 % shorter total duration was observed than for scenario group G12. Accordingly, fleet utilization is increased by 64 % and response times are decreased by 66 %. While it is to be anticipated that all-rounder robots exhibit a better performance, the findings show that the difference is indeed large enough to impede the practical feasibility of specialists fleets. Hence, it can be inferred that the use of all-rounder robots is highly advantageous from a fleet performance perspective. Nonetheless, this poses challenges in robotic development, as integrating diverse functionalities within a single system typically entails increased structural complexity and manufacturing costs.

5.3 Robot Velocity

Since a considerable part of MSR-based task execution is spent moving between different points-of-interest, the driving speed of individual robotic fleet members has a direct impact on the high-level performance of the entire fleet. Using the FleetOR simulation framework, the resulting fleet performances can be compared for different robot speeds, which provides important insights regarding speed levels that are realistically required for meeting the demands of practical applications. Moreover, the merits of potential techniques for reducing driving durations can be evaluated.

Study Design

Since the results presented in the previous section indicate that a robot speed of $v_r = 0.3m/s$ may not be sufficient for practical applications, the merits of introducing zones of different speed within the OR wing was studied (also see section 4.3.3). Such a compromise between speed and safety could allow for achieving acceptable task execution times, while minimizing the risk of collisions. In this regard, the most critical part of the OR wing certainly is the operating room itself, with the sterile zone at its core. Here, keeping driving speeds low can hardly be avoided, since spaces are very narrow, moving objects are omnipresent and hygiene requirements are high. While other parts of the OR wing – in particular hallways and storages – must still be considered challenging, the space and hygiene requirements are somewhat more relaxed and there is more potential for introducing robot-related modifications. Thus, it is argued that higher robot speeds will be acceptable and feasible within these spaces, provided the necessary framework conditions are created.

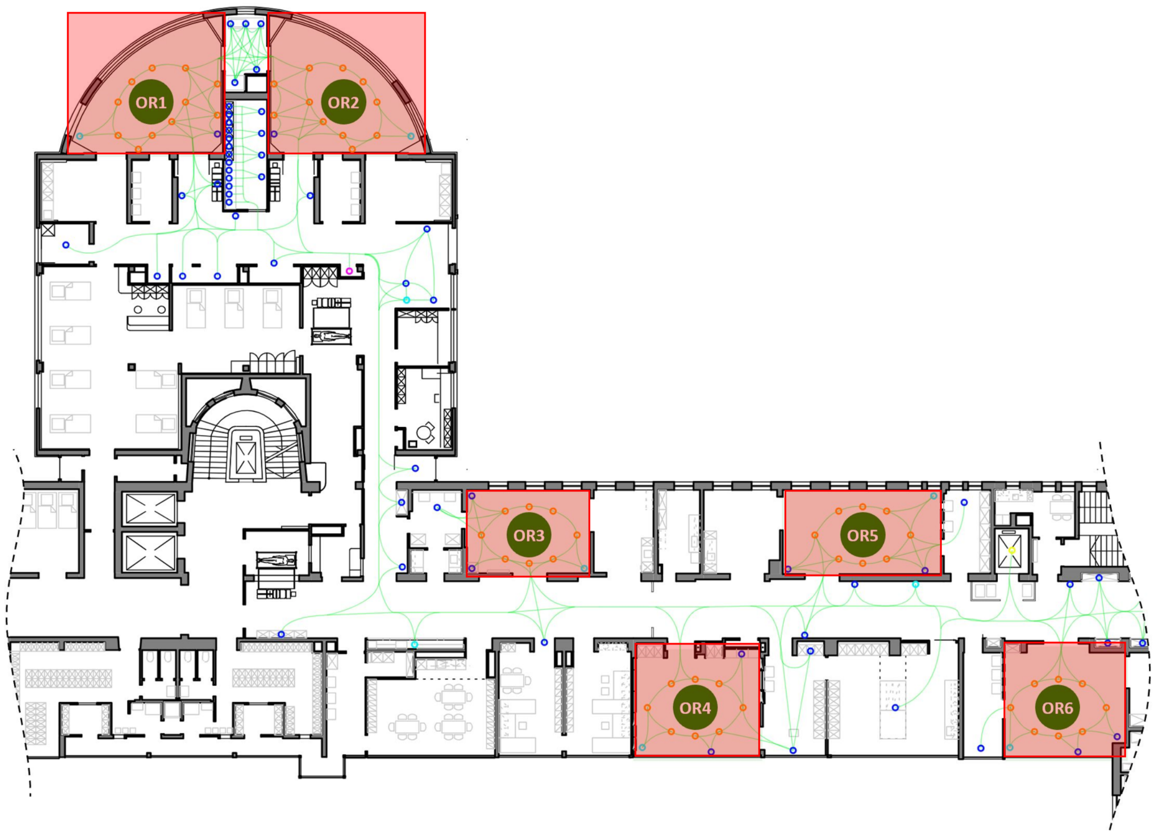

Figure 5.4: This map of the OR wing indicates the speed zones defined for simulation scenario 2. For all zones marked in transparent red, a slow robot driving speed of $v_r = 0.3m/s$ was used, while a fast driving speed of $v_r = 1.2m/s$ was used for all remaining parts of the environment.

To be able to quantify the merits of a hybrid approach regarding robot speed, three scenarios were simulated in FleetOR and compared to each other:

- **Scenario 1:** Robots always move at a slow speed of $v_r = 0.3m/s$.
- **Scenario 2:** Robots move at $v_r = 0.3m/s$ within operating rooms and at $v_r = 1.2m/s$ within all other parts of the OR wing, according to the speed zones defined in Figure 5.4.
- **Scenario 3:** Robots always move at a fast speed of $v_r = 1.2m/s$.

Figure 5.4 shows the speed zones that were defined for the purpose of this study.

Results

Complementing the results presented in section 5.1, the simulation outcomes for the hybrid approach with zones of different speed are shown in Figure 5.5.

The bar plot in figure part 5.5a shows the average task driving time, i.e. the average part of the overall task execution duration that is solely spent for driving between points of interest. As one would expect, scenario 1 with a robot speed of $v_r = 0.3m/s$ resulted in average task driving durations roughly four times longer than those yielded by scenario 3 with a robot speed of $v_r = 1.2m/s$. This corresponds to an increase of 318,6 %. More interestingly, the hybrid scenario 2 (with speed zones) only resulted in an increase in driving time by 77.4 %, which brings the results of the hybrid scenario considerably closer to the fast scenario than to the slow scenario.

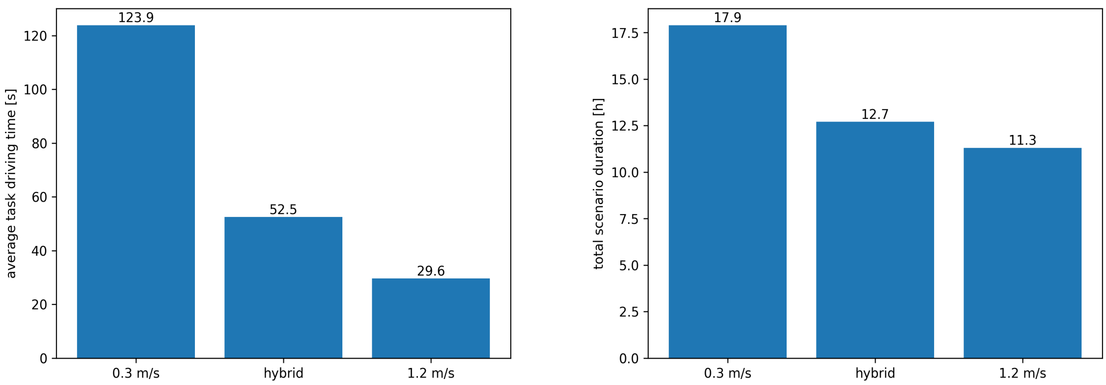

(a) Average task driving time

(b) Total scenario duration

Figure 5.5: The average task driving time (a) and the total scenario duration (b) are shown for simulation scenarios 1-3.

The influence of the different driving speeds can also be observed when considering total scenario durations, as shown in Figure 5.5b. Here, the simulation yielded a decrease of 29,1 % when comparing the hybrid scenario 2 to the slow scenario 1 ($v_r = 0.3m/s$). This is close to the results observed for scenario 3 ($v_r = 1.2m/s$), which yields a decrease of 36,9 %. In essence, the hybrid scenario is much closer to the fast scenario, duration-wise, than to the slow scenario, with the added benefit of slower robot speeds within the operating rooms.

Discussion

Based on the findings presented in the previous, it can be concluded that robot velocity has a strong impact on all studied performance measures. When considering task group G6 (6 ORs, $v_r = 1.2m/s$), for instance, a level of minimum total duration is reached at approximately 10 h, requiring a fleet size of 8 robots or more. In contrast, task group G3 (6 ORs, $v_r = 0.3m/s$) necessitates an increase in fleet size by 75 % to reach its minimum duration level, which is also 27 % higher. Similar trends can be identified for average task response times.

It can further be observed that lower robot velocities result in an increased fleet utilization. This is caused by an overload of the entire system, since slow-moving robots are not able to keep up with demand and pending task requests are accumulating. While a high fleet utilization is typically desirable and not inherently problematic, the observed negative impact on execution times is hardly acceptable for the studied scenarios.

In conclusion, the observations above underscore the pivotal role of robot velocity regarding the feasibility of integrating MSR fleets into the OR wing. When considering key performance parameters of the robotic fleet, as provided by the simulation results, it must be concluded that a slow robot speed of $v_r = 0.3m/s$ is not sufficient for meeting the requirements. First and foremost, the elongation of surgical interventions and, as a consequence, of an entire OR workday is rather extreme. For a scenario with 6 all-rounder robots, 6 ORs and 3 consecutive interventions per OR per day, the conducted simulations yield a total workday duration of approx. 19,5 h for $v_r = 0.3m/s$, which is more than double the recorded human-only performance (9,2 h). This is clearly not acceptable. A robot speed of $v_r = 1.2m/s$, on the other hand, approaches human performance levels quite closely (10,5 h), which is not optimal, but considerably less severe. However, such high robot speeds can hardly be reconciled with the safety demands of operating rooms, which are required due to the presence of moving persons, sterile zones, sensitive equipment etc. Consequently, it must be concluded that both scenarios are not entirely feasible for real-world applications.

As proposed in the previous, a potential resolution to this dilemma involves adopting a hybrid approach, where robots would be required to drive at slower speeds within confined operating rooms while being permitted to navigate less critical areas of the OR wing at higher velocities. While it is obvious that the performance level of the fast scenario cannot be matched using this compromise, the results presented in the previous indicate that the trade-off regarding fleet performance is considerably smaller than for the slow scenario and quite close to the fast scenario.

The implementation of such a hybrid approach, however, might require considerable adaptations to today's OR wing environments. Realizing a robot speed of $v_r = 1.2m/s$ within corridors and other peripheral facilities is still challenging when considering the status quo. This is illustrated by Figure 5.6, which shows that corridors are often used to store medical devices and other equipment that is currently not utilized. In addition, there may be many persons (nurses, doctors, students etc.) moving within the environment, especially when several surgeries are being performed simultaneously. These conditions are hardly optimal for the integration of fast-moving mobile robots.

Following the principles of *surgineering* [Feu+19], it is thus proposed to modify the environment and processes themselves in order to benefit the integration of urgently required technologies. While the options for that are limited for existing hospitals, there are simple measures that can easily be applied without necessitating constructional changes, such as the introduction of clinical regulations governing the storage management within corridors and

Figure 5.6: In the studied real-world OR wing, it is common practice to park equipment – such as medical devices, OR tables, instrumentation tables – in the hallways adjacent to the operating rooms. This often introduces challenging obstacles for mobile robotic systems.

other peripheral OR wing facilities, similar to existing regulations for the industrial domain [DIN20]. By this means, the permanent clearance of robot transit paths could be ensured.

When designing new hospitals [Ama+21], on the other hand, there are many more possibilities for realizing high robot speeds. Similar to the concepts found in industrial environments, "robot-only" zones can be implemented within the architectural layout, which provide a certain separation between human and robotic traffic. There are several concepts with varying degrees of separation (mixed area, restricted area, closed area) and required safety installations (fences, person recognition system, mirrors, markers etc.) [Ver09; DIN20].

5.4 Battery Life

In order to be useful for real-world applications within the OR wing, MSR systems must be continuously available for an acceptable amount of time before having to recharge their batteries. As a result of the requirements analysis conducted for the research project AURORA, it was concluded that the robotic circulating nurse should ideally be available for the duration of an entire surgery without completely emptying its battery in the process. However, the robot may use down-times without pending task requests for moving to the home base to recharge. In the following section, the dimensioning of batteries used within MSRs is investigated in the context of OR wing applications. Based on simulation results, requirements regarding battery specifications are derived and discussed.

Study Design

Using FleetOR, the influence of battery drain rates and charging rates on global fleet performance was studied. Therein, the *drain rate* d refers to the decrease of battery level (ranging from 0.0 to 1.0) per unit of time [1/s], which depends on the activity of the robot and its power demands. For example, driving from point A to point B is associated with a higher drain rate than waiting at a fixed position. On a hardware level, a given drain rate may be the result of different combinations of battery capacity c [Ah] and average battery drain current I [A]:

$$d = \frac{I}{c}$$

Therein, the capacity depends on the choice of battery and the drain current depends on the average power consumption of the robot for a certain activity, which varies among different robot types. For the purposes of the simulations presented in the following, a fleet consisting of 6 robots with identical hardware specifications (and thus identical drain rates) was studied. By defining different alternative simulation scenarios, the drain rates for the basic robot activities *driving* (d_d), *manipulating* (d_m) and *waiting* (d_w) were varied. This effectively changes how long a robot can perform the corresponding activity, before emptying its battery. The maximum durations defined for the studied scenarios are summarized in Table 5.3.

Since the choice of robot speeds heavily influences the resulting driving times (and thus the power consumption), both fast speeds ($v_r = 1.2m/s$, scenarios F1-F16) and slow speeds ($v_r = 0.3m/s$, scenarios S1-S16) were considered. This includes two baseline scenarios with "infinite" battery life ($d_d = d_m = d_w = 0$) for each of both robot speeds (scenarios F1 and S1), which were introduced to determine the fleet's performance without the limitation of battery life.

As explained in section 4.3.5, the battery simulation triggers a forced recharge if the battery level of a robot falls below a threshold b_{low}. In this case, the robot is ordered to drive back to its base and charge at least until a defined battery level b_{high} is reached. For the sake of the presented simulations, the lower threshold was set to $b_{low} = 0.1$ and the upper threshold to $b_{high} = 0.5$, aiming at a good balance between up and down times. The recharge rate was chosen such that a full charge can be achieved in 30 minutes (2C rating), which is in line with the capabilities of currently available batteries for AGV-style systems.

Furthermore, it is important to note that a robot is sent back to its base by the fleet manager in situations without any pending task requests. Naturally, the time spent at the base is used for charging the robot, however, this process can be interrupted at any time, in contrast to a forced recharge. In the following, only forced recharges are considered, since these result in a temporary downtime of the robotic resource and thus should be avoided. In the optimal case, a robot is able to sustain its battery level without recharging or only by means of intermediate charges.

Results

Figure 5.7 shows the total amount of forced recharges across the entire fleet for all considered scenarios. For the fast scenarios (F1-F16), no recharges were required up to F7, which means that a robot with enough battery capacity to drive for 50 minutes or more is able to sustain its

Table 5.3: For the simulated fast and slow scenarios (F1-F16 and S1-S16), the defined maximum durations for driving, manipulating and waiting that can be achieved with a fully charged battery, are given. For example, the drain rates for scenario F6 were chosen such that the robot is either able to drive for 1 hour, or manipulate for 2 hours, or wait for 4 hours when starting with a full battery.

scenario ID	fast ($v_r = 1.2 m/s$)	F1	F2	F3	F4	F5	F6	F7	F8
	slow ($v_r = 0.3 m/s$)	S1	S2	S3	S4	S5	S6	S7	S8
maximum duration [min]	driving	∞	300	150	100	75	60	50	43
	manipulating	∞	600	300	200	150	120	100	86
	waiting	∞	1200	600	400	300	240	200	171
scenario ID	fast ($v_r = 1.2 m/s$)	F9	F10	F11	F12	F13	F14	F15	F16
	slow ($v_r = 0.3 m/s$)	S9	S10	S11	S12	S13	S14	S15	S16
maximum duration [min]	driving	38	33	30	27	25	23	21	20
	manipulating	75	67	60	55	50	46	43	40
	waiting	150	133	120	109	100	92	86	80

battery level solely by intermediate charging. For lower values, the number of required forced recharges increases in a linear fashion, passing a count of 38 recharges for a maximum driving duration of 20 minutes.

For the slow scenarios (S1-S16), a maximum driving duration of 300 minutes or more is required to avoid forced recharges. Again, the number of recharges increases linearly for decreasing driving durations. However, the slope of this increase is considerably steeper than for the fast scenarios. For scenario S16, with a maximum driving duration of 20 minutes, the number of required recharges is 336.

The influence of forced recharges on overall fleet performance regarding total scenario duration is shown in Figure 5.8. For the fast scenarios, the impact is comparatively small, with an added duration of 6.47 % (approx. 44 min) for scenario F16. For the slow scenarios, the impact is quite extreme, with an added duration of 71.75 % (approx. 12 h and 16 min) for scenario S16.

Discussion

In the OR context, a large part of task requests is performed on demand, which aggravates the planned charging of robotic systems. While the course of a surgical intervention usually follows an established workflow, the exact progress is often subject to various influences and, thus, is hard to predict accurately. There may be situations with ongoing high demand and no opportunity for intermediate recharges followed by phases with a comparatively low workload where robots spend more time at their home bases. It is therefore desirable to dimension the power supply of robots such that phases of high demand – where forced recharges are especially obstructive – can be bridged effectively. On the other hand, it is important to use spontaneous declines in workload for intermediate charges, since battery capacities cannot be increased indefinitely due to the size and weight restrictions of mobile systems.

The simulation results indicate that such a compromise can rather easily be found for fast-

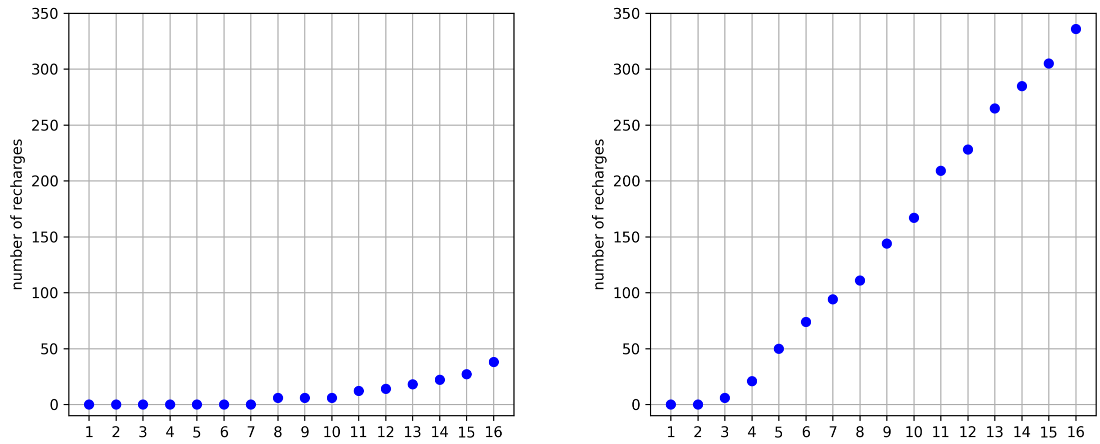

(a) fast scenarios (F) **(b)** slow scenarios (S)

Figure 5.7: The number of forced recharges across the entire fleet is shown for all simulated scenarios. Figure part (a) gives results for the fast scenarios (F1-F16), while part (b) is concerned with the slow scenarios (S1-S16).

moving robots: A fleet of 6 robots with battery capacities for achieving a maximum driving duration of up to 50 min (F7), is able to serve 6 ORs without requiring forced recharges and thus without hampering overall fleet performance compared to the reference scenario F1. When considering the capabilities of today's battery systems for mobile robots, the realization of such a configuration is reasonable.

For slow-moving robots, the capabilities of the installed battery solutions is considerably more critical. To not hamper overall fleet performance, a battery capacity supporting driving durations of 5 h a would be required to avoid forced recharges for the investigated scenarios. While not being entirely unreasonable, achieving such high performance levels requires large battery systems, which complicates the design of slim and light-weight robotic systems.

5.5 Task Prioritization

As motivated in section 4.4.2, the prioritization of tasks is a central part of MSR fleet management within the OR wing. Two prioritization strategies, FIFO and ANTS-OR were described, which are compared in the following.

Study Design

Using the FleetOR simulation framework, the two prioritization approaches were evaluated for the same scenario in order to identify differences in fleet performance. A fleet of 4 all-rounder robots ($v_r = 1.2m/s$) was configured to serve 6 operating rooms. By design and with reference to the results presented in section 5.1, this creates a constellation where demand exceeds supply, i.e. the robotic fleet is likely to face situations where not enough robotic resources are available to instantly execute all currently requested tasks. In these situations, the ANTS-OR approach ANTS-OR is expected to achieve a more differentiated task prioritization than a simple FIFO approach. In each OR, three consecutive interventions were simulated.

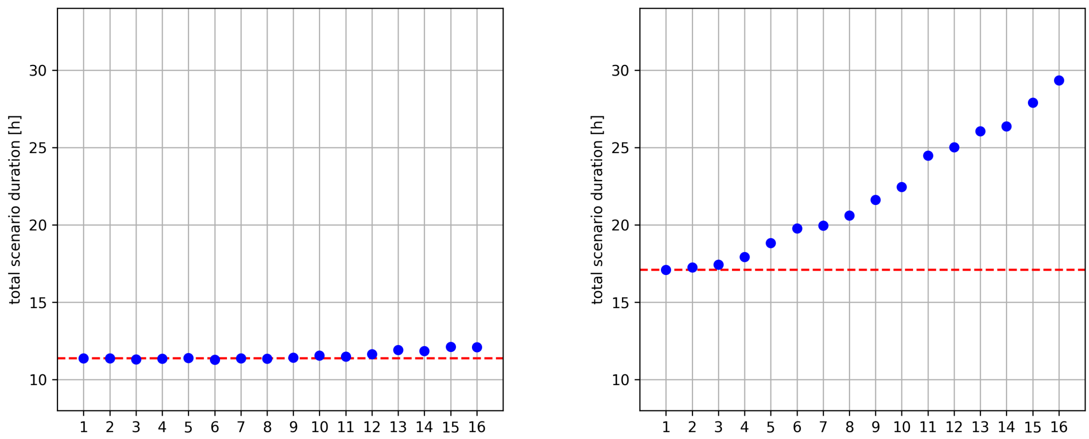

(a) fast scenarios (F) **(b)** slow scenarios (S)

Figure 5.8: The total scenario duration is shown for all simulated scenarios. Figure part (a) gives results for the fast scenarios (F1-F16), while part (b) is concerned with the slow scenarios (S1-S16). For better readability, the results obtained for the reference scenarios with infinite battery capacity (F1 and S1) is additionally highlighted by dashed red lines.

Results

A high prioritization of critical tasks is expected to decrease their task response time, i.e. the duration from task request to start of execution, which is also referred to as idle duration. Figure 5.9 summarizes the idle durations yielded by FIFO and ANTS-OR for tasks with different emergency levels. The distributions are depicted as box plots, with the median, upper quartile and lower quartile marked by horizontal lines. When comparing only FIFO distributions to each other using the Kruskal-Wallis test, no significant difference can be shown between the different emergency levels ($p = 0.804$). This is to be expected, since FIFO prioritization does not consider emergency levels. For ANTS-OR, on the other hand, a significant difference among the two groups can be shown ($p < 0.001$).

When comparing FIFO to ANTS-OR, it can be observed that the idle durations of emergency level E-2 are shorter for ANTS-OR than for FIFO, with a 74.2 % lower median. Since distributions are non-Gaussian, a Mann-Whitney U test was used to confirm that the groups show a statistically significant difference ($p < 0.001$, $\alpha = 0.05$). As expected, the preference of E-2 tasks is at the expense of tasks with emergency level E-0, where longer idle durations can be observed compared to FIFO prioritization. Here, the median shows an increase of 152.8 %, which, again, is a statistically significant difference according to the Mann-Whitney U test ($p < 0.001$). For an emergency level of E-1, both prioritization strategies yield quite similar results with a 4.3 % lower median for ANTS-OR prioritization. Still, this difference is significant ($p = 0.012$). Complete numerical results for mean, median and mean absolute deviation are given by Table 5.4.

Figure 5.10 shows idle durations for all different command levels. As expected, the distributions for FIFO do not show a significant difference when compared to each other using the Kruskal-Wallis test ($p = 0.488$). For the ANTS-OR case, however, the idle durations are short for the highest-priority command level D-3 (mean: 320.00 s, median: 55.17 s) and increase continuously until reaching the two lowest-priority levels D-0 (mean: 630.32 s, me-

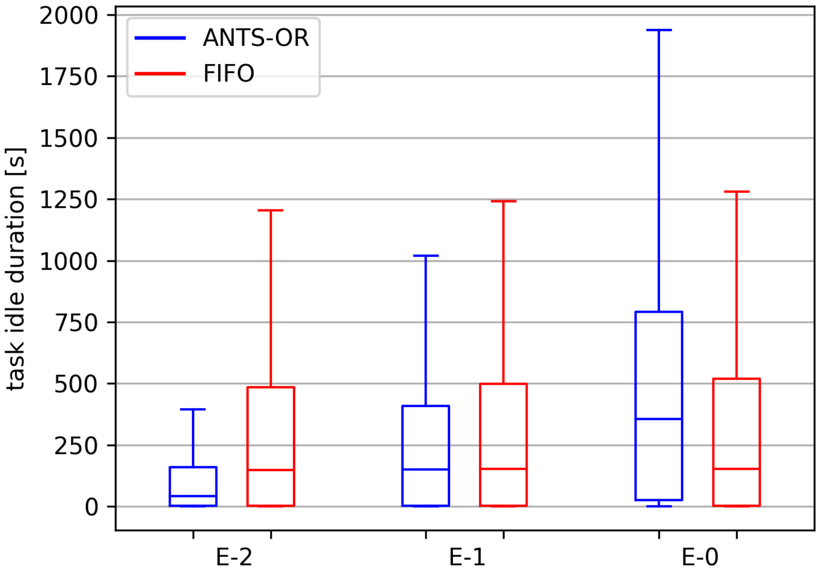

Figure 5.9: Task idle durations for the different emergency levels are shown. As expected, FIFO is not influenced by the emergency level. On the contrary, the idle duration of high-priority E-2 tasks is significantly reduced by ANTS-OR prioritization, at the expense of low-priority E-0 tasks.

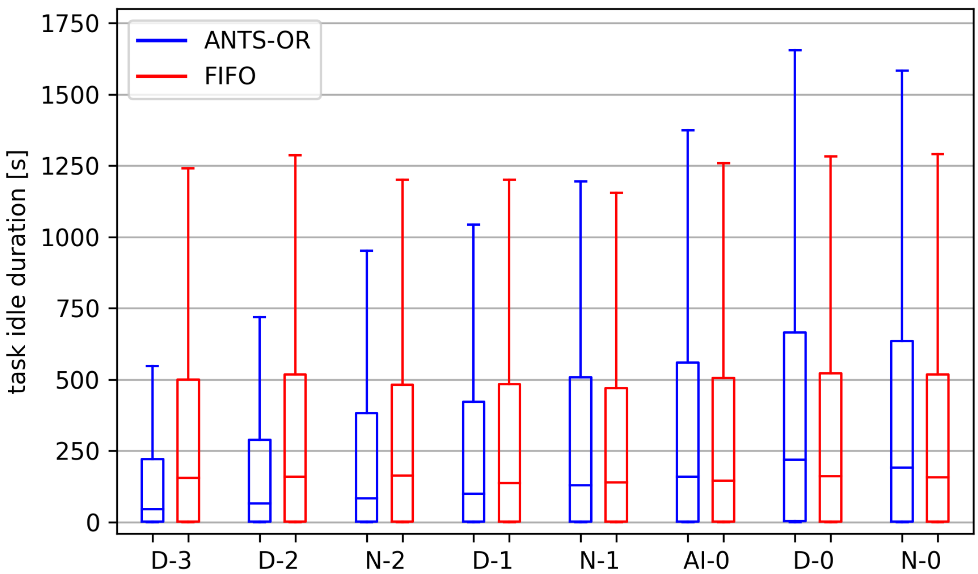

Figure 5.10: Task idle durations for the different command levels are shown. As expected, FIFO is not influenced by the command level. On the contrary, the idle duration of high-priority tasks (D-3, D-2, N-2, D-1) is significantly reduced by ANTS-OR prioritization, at the expense of low-priority tasks (AI-0, D-0, N-0).

Table 5.4: Task idle durations are shown for both the ANTS-OR and the FIFO scenario. For each emergency and command level, the mean and median duration is given, along with the mean absolute deviation (MAD). The p-value indicates whether ANTS-OR and FIFO distributions are significantly different.

Level	ANTS-OR			FIFO			p
	mean [s]	median [s]	MAD [s]	mean [s]	median [s]	MAD [s]	
E-2	168.99	43.67	200.59	490.58	169.09	537.03	<0.001
E-1	429.49	152.36	483.29	472.44	146.12	523.09	0.012
E-0	832.59	370.32	889.12	470.07	146.43	526.40	<0.001
D-3	320.00	55.17	407.03	447.71	145.06	494.92	<0.001
D-2	343.92	73.55	410.75	493.65	155.71	545.77	<0.001
N-2	449.86	95.8	535.09	466.59	140.73	523.45	0.024
D-1	437.63	96.02	526.32	521.6	171.08	580.23	<0.001
N-1	469.49	114.91	550.26	466.36	139.52	520.46	0.169
AI-0	565.77	141.69	654.34	472.59	156.79	516.29	0.242
D-0	630.32	178.37	717.11	473.78	165.82	514.66	0.004
N-0	603.12	174.29	683.37	477.18	151.76	531.7	<0.001

dian: 178.37 s) and N-0 (mean: 603.12 s, median: 174.29 s). Using the Kruskal-Wallis test, a significant difference of the command levels can be shown ($p < 0.001$).

By means of the Mann-Whitney U test, idle times of a given command level can be compared between FIFO and ANTS-OR. A statistically significant difference can be shown for the highest levels (D-3, D-2, N-1, D-1), where ANTS-OR yields shorter idle durations than FIFO. Accordingly, a statistically significant difference can be shown for the lowest levels (D-0, N-0), where ANTS-OR yields longer idle durations. For the remaining intermediate levels (N-1, AI-0), no significant difference can be shown, which indicates that similar idle durations can be expected for both prioritization strategies. When comparing the idle times yielded by ANTS-OR for D-0 to those for N-0, no significant difference can be shown using the Mann-Whitney U test ($p = 0.083$). This is to be expected, since both levels are prioritized equivalently (also see Figure 4.8). Complete numerical results for mean, median, mean absolute deviation and p-values are given in Table 5.4.

For the ANTS-OR scenario, a slight increase of total scenario duration by 1.46 % (10 min, 10 s) can be observed, compared to FIFO.

Discussion

Based on the results presented in the previous, it can be concluded that ANTS-OR prioritization effectively modulates the idle duration of tasks according to the current situation. Tasks that are related to an emergency situation with critical patient condition are executed with preference, leading to shorter waiting times. On average, an emergency task (level E-2) is executed after 2 min 49 s, which is 5 min 22 s earlier compared to FIFO scheduling. This is a considerable time difference in the context of critical situations. Using ANTS-OR, urgently needed materials can be provided sooner and medical device adjustments can be completed within a shorter timespan. At the same time, less critical tasks related to routine situations are postponed in case other more important task requests are pending. Compared to FIFO prioritization, this results in longer idle durations and thus longer waiting times for these tasks, which is a trade-off that must be taken into account when using ANTS-OR. Regarding command levels, similar results were obtained. The task requests of higher-ranking team members are served with priority over those originating from lower-ranking team members or digital information systems. Overall, it can be concluded that the intended effect of ANTS-OR was demonstrated by the simulation results. Using ANTS-OR, a more differentiated and situation-dependent task prioritization was achieved, which makes it more suitable for OR wing workflows.

An important limitation of ANTS-OR prioritization is that emergency-related tasks are not further compared regarding urgency. For example, there might be several emergency patients at the same time with different conditions and needs. This is not considered by ANTS-OR and may involve the resolution of ethical predicaments, as is further explored in section 6.3. Furthermore, due to the introduction of further prioritization criteria beyond idle duration (namely emergency and command levels), a certain increase in total scenario duration must be taken into account for ANTS-OR. However, with an elongation of just over 10 minutes for an entire workday of interventions, the impact of this effect is minimal and should not make a noticeable difference in practice.

5.6 Mission Planning

In section 4.4.3, the *Vehicle Routing Problem for the OR Wing* (VRP-OR) was stated and a greedy heuristic solution algorithm was proposed. In the following, a simulation-based study is presented, aiming at evaluating this novel task allocation approach and quantifying its benefits and drawbacks with regard to global performance parameters of the robotic fleet. Parts of this study were published in [BKW23].

Study Design

Using the FleetOR simulation framework, a virtual robotic fleet consisting of 6 all-rounder robots ($v_r = 1.2m/s$) serving 6 operating rooms was studied. This is in line with the recommended fleet size for this number of ORs, according to the results presented in section 5.1. Three consecutive surgical interventions were scheduled for each OR, representing a typical work day. A set of 8 alternative scenarios were simulated, resulting from varying the maximum length of robotic missions (i.e., the maximum number of requests per mission) from 1 to 8. Of note, a maximum mission length of 1 results in the trivial case of single-task scheduling with instantaneous assignment, which was used as a reference for comparison. Since OR wing workflows are highly dynamic and prone to frequent changes, a mission length greater than 8 was not considered. It is assumed that each robot possesses an inventory storage that is large enough to fit all objects required for executing the tasks of a given mission, which is reasonable for the studied mission lengths.

As described in section 4.4.3, a greedy routing strategy was proposed for addressing the alternative pickup locations that are intrinsic to the VRP-OR problem[1]. The performance of this approach was evaluated and compared to exhaustively searching for the optimal route. In particular, the differences in fleet performance and computational effort were studied.

Results

Figure 5.11a shows resulting task driving durations for all studied mission lengths. A mean value of 29.29 s, a median of 25.82 s, and a mean absolute deviation of 18.37 s were observed for the reference scenario (mission length 1). When increasing the mission length, driving durations decrease in an inverse exponential behavior. This leads to a mean value of 15.02 s and a median of 8.08 s for a mission length of 4, achieving a decrease by 48.7 % and 68.71 %, when comparing to the reference scenario. Likewise, a mean value of 11.41 s and a median value of 3.48 s were observed for a mission length of 8, leading to decreases by 61.04 % and 86.54 %. According to the Mann-Whitney U test, each distribution shows a statistically significant difference to the reference scenario ($p < 0.001$). Complete numerical results for mean, median, mean absolute deviation, and percentual relative improvements are given in Table 5.5.

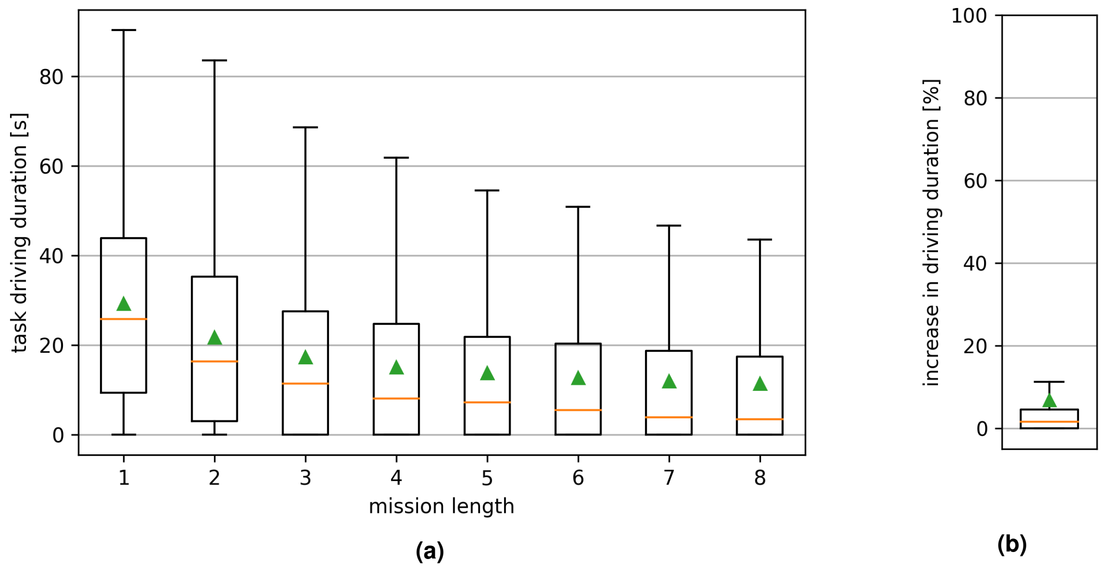

Figure 5.11: Figure part (a) shows the (single-)task driving durations that were obtained for different mission lengths, when using the proposed VRP-OR solution approach. Figure part (b) shows the percentual increase in driving time when comparing the results of the greedy algorithm to an optimal routing strategy. Median values are indicated by orange lines, mean values are depicted as green triangles.

The performance of the greedy routing algorithm compared to exhaustive searching is shown in Figure 5.11b, for missions of length 8. When using greedy routing, an increase in driving duration by 6.78 % (mean) and 1.56 % (median) must be taken into account, compared to an optimal solution strategy. However, a maximum computation duration of

[1] It must be emphasized that this explicitly refers to the routing step within the mission planning algorithm described in section 4.4.3. While the overall algorithm is also based on the greedy paradigm, a comparison to the optimal case cannot be given due to impossibly long computation times.

Table 5.5: The achieved driving durations are shown for the studied mission lengths. For each scenario, the mean and median driving durations are given, along with the relative decrease when comparing to the reference scenario (mission length 1). Furthermore, the mean absolute deviation (MAD) describes the variance of each distribution and the p-value indicates whether the observed difference to the reference scenario is statistically significant.

mission length	mean [s]	mean decrease [%]	median [s]	median decrease [%]	MAD [s]	p
1	29.29	-	25.82	-	18.37	-
2	21.76	25.71	16.3	36.88	17.51	< 0.001
3	17.27	41.04	11.36	56.01	15.46	< 0.001
4	15.02	48.7	8.08	68.71	14.18	< 0.001
5	13.78	52.95	7.17	72.23	13.44	< 0.001
6	12.74	56.49	5.51	78.66	12.86	< 0.001
7	11.97	59.14	3.92	84.82	12.57	< 0.001
8	11.41	61.04	3.48	86.54	12.31	< 0.001

2.48 h was observed for optimal routing, when performing a full scheduling cycle. For greedy-based routing, on the other hand, a maximum computation duration of only 38.9 ms was observed.

Discussion

The results presented in the previous demonstrate that the proposed algorithm is indeed well-suited for solving the VRP-OR problem. Compared to the baseline of single-task allocation, the proposed multi-task mission planning method allows for the elimination of redundant trips by combining task requests that can be jointly executed along an optimized route. Even though slightly longer driving durations were observed for the greedy routing approach, these are more than acceptable when considering the extremely short processing times, which cannot be achieved with optimal solvers due to the underlying optimization problem being NP-hard. The maximum scheduling time that was observed during the simulations is well below the acceptable threshold defined in the requirements section 4.4.1, meaning that the fleet managing system is able to quickly react to changing circumstances.

It can also be concluded that increasing the mission length indefinitely is not reasonable, since an asymptotic behavior can be observed for the improvements in idle and driving durations. When considering the hard-to-predict workflows taking place within the OR wing, keeping mission lengths rather short seems reasonable as well, since long missions would be subject to frequent rescheduling anyway. Based on the results, a mission length of 5 or 6 can be recommended, since these values provide a good compromise between performance and mission length. For longer missions, performance gains increase only slowly.

5.7 Quick Storages

By introducing additional storages – from here on referred to as *quick storages* – within the environment, objects that will likely be needed in the near future can be prepared in advance and, upon request, provided to the surgical team within a shorter amount of time. Since realistically achievable robotic driving speeds are very limited within the OR wing setting, such approaches can help reduce the performance gap to human circulators. In the following study, the merits and limitations of such a proactive fleet behavior are investigated.

Study Design

In order to proactively and context-dependently fill the quick storages, surgical materials that are likely to be needed in the immediate future must be predicted. For the purposes of this study, this is simulated using the workflow module of FleetOR, which is aware of the next upcoming task request for each OR based on the workflow recordings[2]. Herein, the simulated accuracy of the prediction can be freely adjusted in order to study the impact of incorrect predictions on the overall performance.

One quick storage was introduced within each operating room (see Figure 5.12), located within the non-sterile zone of the OR right next to the position where sterile material is handed over to the surgical team by the robots. In case a given material is requested and it happens to be available from the quick storage, the robot can move directly to the hand-over position, retrieve the material from the quick storage and hand it over. Thus, there is no need to first navigate to the normal storage and collect the material from there, which is expected to reduce the average driving duration significantly.

However, the inventory of the quick storages must be continuously updated, which causes additional workload for the robotic fleet. Firstly, predicted articles must be collected from the normal storage and brought to the quick storages. Secondly, articles that were erroneously predicted and are not actually needed for the surgical intervention must be brought back to the main storage in order to free space for the next articles and to avoid an overflow of the quick storage. In order to fit well within the OR environment, a quick storage should be small and thus only able to hold a very limited amount of articles. To simulate this limitation, a threshold amount was introduced, at which the fleet manager starts to generate task requests for returning surplus articles. This threshold can be adjusted freely to study the impact of different quick storage capacities. For the following study, a threshold of 2 articles was used, which should be realistically achievable when considering the space situation within the OR. The article return process follows a first-in-first-out scheme, meaning that surplus articles are returned in the order they were delivered. Task requests for stocking the quick storages and returning surplus articles are generated by the fleet manager and assigned with a low priority. This is to ensure that more urgent tasks that were directly requested by the surgical team are executed with preference, effectively using temporary "down-times" – where no user-originating task requests are pending – for executing less time-critical preparation tasks. As long as the overall fleet workload is sufficiently low, this approach is expected to alleviate the negative impact of additional overhead caused by executing preparation tasks.

[2] While the development of a task prediction algorithm was not within the scope of this thesis, first advances have been made in the framework of the AURORA project, most notably by building a knowledge graph which models the demand of surgical materials in different situations [Mül+20].

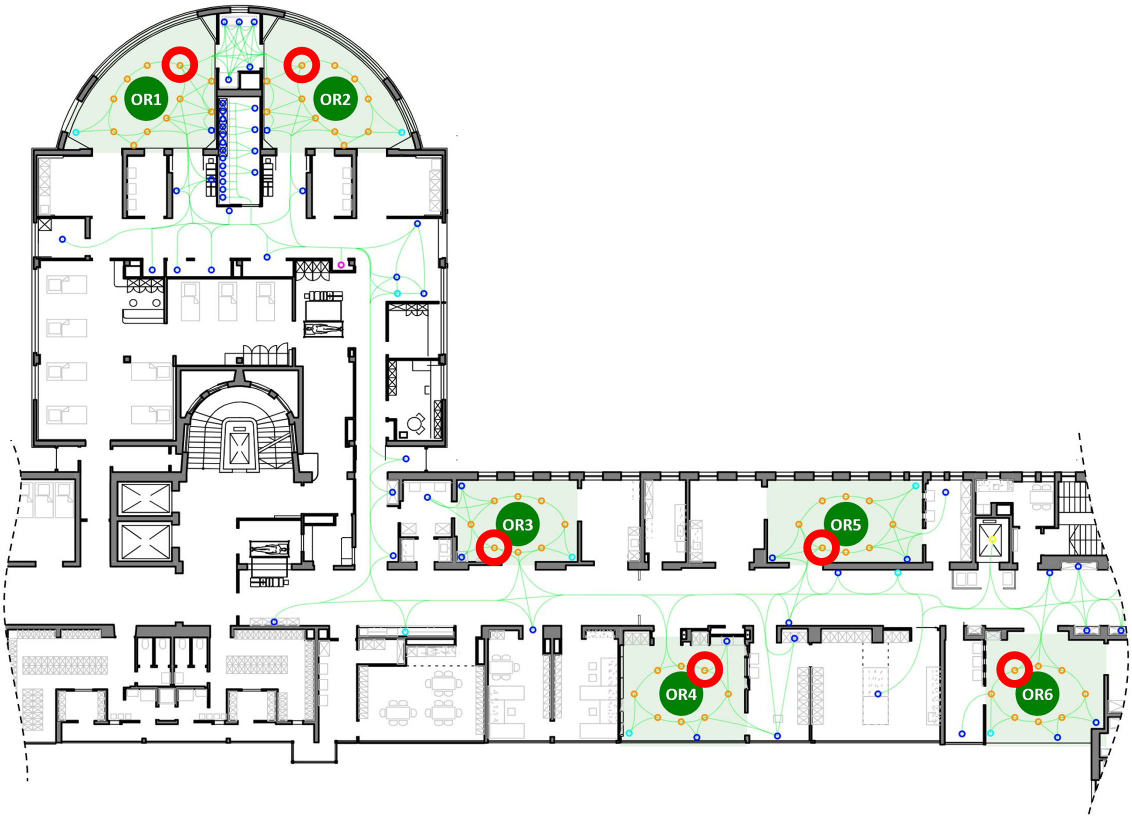

Figure 5.12: This map of the OR wing indicates the locations of all quick storages that were introduced for the purposes of this study.

Results

Figure 5.13 shows the impact of the introduced quick storages on task driving times. This only relates to tasks of type "T1" (transportation of light materials), since only these tasks involve materials that can realistically be stored within quick storages. Furthermore, only intraoperatively requested tasks are considered; preoperative preparation tasks do not benefit from a task prediction, since they are not performed on demand and the list of articles to be prepared is already known in advance, e.g., from clinical guidelines. As a baseline for the comparison, results for a reference scenario ("ref") are given as well, which indicates the performance without any proactive behavior.

As expected, the mean and median driving durations decrease with increasing prediction accuracy. While rather modest improvements of 7.37 % (mean) and 14.71 % (median) can be observed for a low accuracy of 0.2 (see Table 5.6), a considerable improvement is achieved for higher accuracies: For 0.8, the mean is decreased by 25.20 % and the median by 43.04 % with regard to the reference scenario. A hypothetical "perfect" prediction algorithm (accuracy of 1.0) leads to improvements of 32.11 % (mean) and 45.95 % (median), which thus represents the upper threshold in performance that can be achieved by introducing quick storages. According to the Mann-Whitney U test, the improvement of each scenario is statistically significant ($p < 0.001$) when compared to its respective predecessor of lower accuracy.

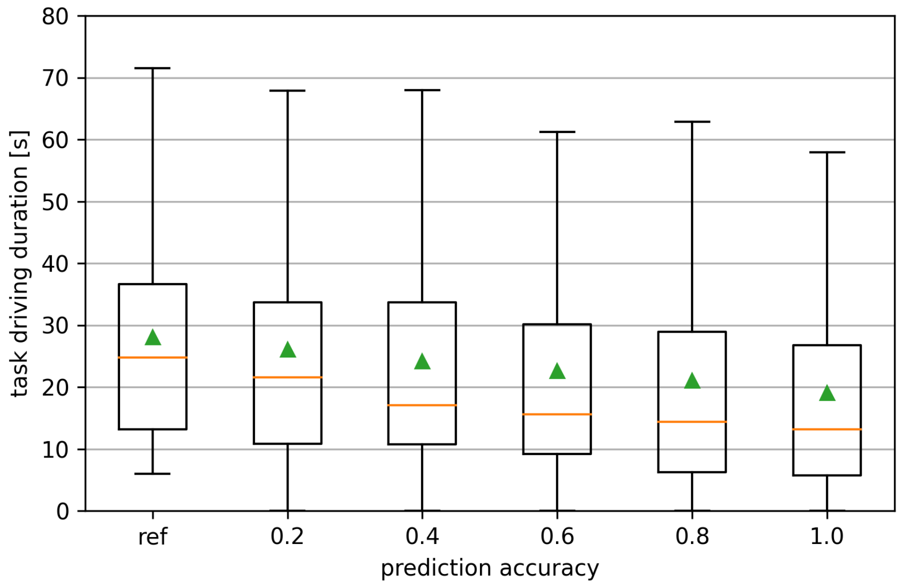

Figure 5.13: Driving durations of intraoperative T1 type tasks are shown as boxplots for scenarios with article prediction accuracies ranging from 0.2 to 1.0. As a baseline for comparison, results for a reference scenario (ref) without proactive behavior are given as well. Median driving durations are indicated by orange lines, mean driving durations are depicted as green triangles.

Table 5.6: Numeric results for driving durations of T1 type tasks are given for different task prediction accuracies. For each simulated scenario, the mean and median driving durations are provided, as well as the relative percentual decrease when comparing both measures to the reference scenario. Furthermore, the mean absolute deviation (MAD) is given as a measure for the variance of each distribution. The p-value provided in the last column is the result of the Mann-Whitney U test, when comparing each scenario to its predecessor.

prediction accuracy	mean [s]	mean decrease [%]	median [s]	median decrease [%]	MAD [s]	p
ref	28.09	-	24.72	-	14.44	-
0.2	26.02	7.37	21.55	14.71	14.71	< 0.001
0.4	24.15	14.03	17.10	30.83	15.15	< 0.001
0.6	22.59	19.58	15.54	37.14	15.02	< 0.001
0.8	21.01	25.20	14.08	43.04	14.93	< 0.001
1.0	19.07	32.11	13.36	45.95	14.48	< 0.001

To demonstrate that the additional workload introduced by returning falsely predicted articles to the normal storage does not unacceptably affect the performance of other tasks, Figure 5.14 depicts the distributions of total task durations (time of request to time of completion) of all other task types. In case of an adverse effect due to the additional execution of return tasks, these durations would be elongated compared to the reference scenario. Such an impact was not observed, which is further confirmed by the Kruskal-Wallis test, showing no significant difference between the scenarios ($p = 0.470$).

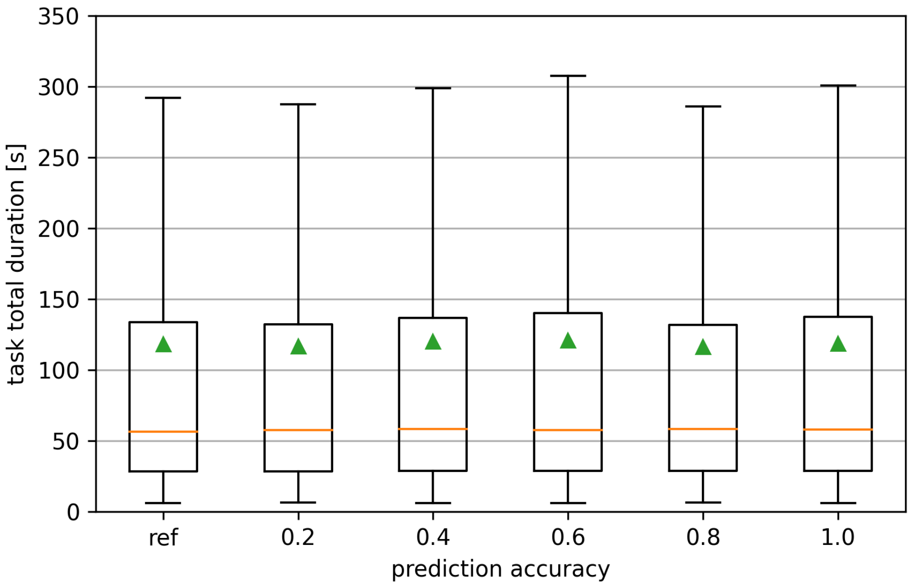

Figure 5.14: Execution durations of all intraoperative tasks (except T1 type tasks) are shown as boxplots for all simulated scenarios. Median execution durations are indicated by orange lines, mean execution durations are depicted as green triangles.

Discussion

The results presented in the previous clearly demonstrate the added benefit of the quick storage approach. Task driving durations are significantly reduced, while the overhead generated by returning unused articles does not noticeably impact overall fleet performance. As demonstrated, it is both beneficial and feasible to use times of low workload for executing preparation tasks.

For using quick storages to their fullest potential, however, a high-accuracy prediction algorithm is required. With an algorithm of perfect accuracy, median driving durations for delivering requested materials could be almost reduced by half, which – from the perspective of the surgical team – results in a considerable decrease of waiting time. Thereby, the gap to human performance levels can be narrowed, without having to increase robotic driving speeds. While a perfect prediction of articles is hardly achievable, scenarios with lower prediction accuracies have still yielded usable results when compared to the reference scenario: A prediction accuracy of 0.8 allows for approaching the optimal performance level quite closely and an accuracy of 0.6 still allows for decreasing median driving durations by over a third. Such improvements would certainly make a welcome and noticeable difference in real-world workflows.

5.8 Robot Inventory

Depending on the robotic design, a storage can be integrated into the mobile base of the robot itself. Such a *robot inventory* is limited by the dimensions of the robot and the sizes of the articles to be stored. As a result of the research project AURORA, a robot inventory of 10, or even up to 20 articles is realistically achievable, at least for smaller sterile goods, such as suturing material and surgical gloves. This raises the question of how a robot inventory can be optimally used to benefit the overall fleet performance. In particular, two approaches were identified:

1. An AI algorithm is used to predict sterile articles that are likely needed in the immediate future. A robot can then flexibly stock its inventory according to the current result of the prediction. This involves the return of articles that were falsely predicted. This approach is best suited for stand-alone robots that are focused on a single OR and are not part of a globally managed fleet. When considering a fleet of robots that is flexibly serving multiple ORs, a predictive approach is less promising, since the requirements of each OR may differ widely at a given point in time, depending on the type and current phase of the respective surgical procedures. Reflecting all these needs simultaneously within a limited robot inventory seems hardly realistic and would lead to a high volatility of the currently optimal inventory stocking, in turn resulting in an elongation of driving durations.
2. Based on statistical data, each robot can be equipped with the most often used articles, irrespective of the current phases of simultaneously performed surgical procedures. For example, a robot with an inventory size of 10 could hold one unit of each of the 10 most commonly used materials. After one of these materials has been requested and removed from the inventory, it must be restocked at a normal storage, which ideally is done during times of low workload.

Since a global management and optimization of the robotic fleet is aimed at in the context of this doctoral thesis, the merits and limitations of the second approach are evaluated in the following.

Study Design

As a prerequisite for the stocking of robot inventories, a statistical analysis of the recorded interventions (see section 4.2) was conducted. Table 5.7 lists the 10 most often used article types that are typically requested during the intraoperative part of laparoscopic cholecystectomies. Thus, from a statistical perspective, it is sensible to hold these articles readily available within robot inventories. Based on these results, multiple scenarios were compared with each other using the FleetOR simulation framework:

- As a baseline for comparison, a reference scenario without robot inventories was simulated.
- Secondly, scenarios with robot inventories containing the 5 most often requested article types were studied. Herein, the number of items per article type was varied between 1 and 4. This results in robot inventory sizes of 5, 10, 15 and 20 articles. As a reference

and theoretical upper performance limit of robot inventories, an additional scenario with an infinite number of items per article type was simulated (inf).

- Similarly, scenarios with robot inventories containing the 10 most often requested article types were studied, resulting in robot inventory sizes of 10, 20, 30 and 40 articles. Again, a scenario with infinite supply per article type was simulated (inf).
- Lastly, inventories containing the 20 most often requested articles were studied, with robot inventory sizes of 20 and 40 articles and a reference scenario with infinite supply (inf).

For each scenario, FleetOR was configured with a fleet of 6 all-rounder robots ($v_r = 1.2 m/s$) serving 6 operating rooms. This is in line with the recommended fleet size for this number of ORs, according to the results presented in section 6.1. Within each operating room, three consecutive surgical interventions were simulated.

In the following, robot inventory sizes are denoted in the form [number of article types] x [number of items per article type]. For example, an inventory size of 5 x 2 refers to an inventory providing the 5 most commonly used article types and stocking 2 items per type (i.e., 10 items in total).

Table 5.7: A list of the most often requested articles during the intraoperative part of laparoscopic cholecystectomies.

rank	name	description
1	Prolene 0 FSL 45cm EH7920H	suturing material
2	Prolene 3/0 FS2 45cm 8665H	suturing material
3	Monocryl 3/0 FS2 45cm Y293H	suturing material
4	WUNDVERBAND CURAPOR 7X5 CM 32912	wound dressings
5	Handschuh OP Latex Biogel Gr.7,5 82275	surgical gloves
6	Easyflowdrainage 6mmx30cm 97.816.92.135	wound drainage
7	Robinson-Drainage Ch. 21	wound drainage
8	Ligaclip Magazin gelb LT400	laparoscopic clips
9	OP-Mantel SMMS unverstärk 150cm XL 19362	surgical gown
10	Vicryl 0 70cm Vio geflochten V987H	suturing material

Results

As is the case for the quick storage approach (see section 5.7), the introduction of robot inventories impacts those tasks for which a given robot would normally have to navigate to a storage to collect a requested article. As defined in section 4.3.4, this type of tasks is denoted as "T1". By means of robot inventories, trips to the storage can be avoided for T1 type tasks in cases where the requested article happens to be present within the inventory. This effectively reduces the driving duration of these tasks.

Accordingly, Figure 5.15, Figure 5.16 and Figure 5.17 show the effect of different inventory sizes on the driving duration of T1 type tasks. As can be seen from the plots, inventories with 5 articles only have a rather small impact on the driving durations, even for scenarios with a higher number of items available for each article type. Even when considering the hypothetical scenario with an infinite supply per article type, the mean driving duration is only decreased by 7.36 % and the median driving duration only by 11.51 % (see Table 5.8), when comparing to the reference scenario without any robot inventories.

However, for inventories providing 10 different article types, driving durations can be reduced more substantially. Notably, peak performance is already reached for a robot inventory size of 10 x 2, which results in a mean decrease of 15.76 % and a median decrease of 35.16 % when comparing to the reference scenario without any robot inventories. Increasing the number of items per article type, including the theoretical scenario with an infinite supply, results in a very similar performance level without a statistically significant difference, which is confirmed by the Mann Whitney U test (see p-values in Table 5.8).

When considering robot storages equipped with 20 different article types, the performance can be further increased, however, only by a small margin and only for robot inventories of 20 x 2 or larger. Again, it can be observed that stocking more than two items per article type does not result in a significantly better performance.

Similar to the quick storage approach (see section 5.7), the additional workload caused by restocking the robot inventories intraoperatively does not significantly impact the performance of other tasks, e.g. by increasing their total duration (request to completion). This is confirmed by the Kruskal-Wallis test, which shows no significant difference between the scenarios ($p = 0.419$ for 5 article types, $p = 0.391$ for 10 article types and $p = 0.842$ for 20 article types, respectively including the reference scenario).

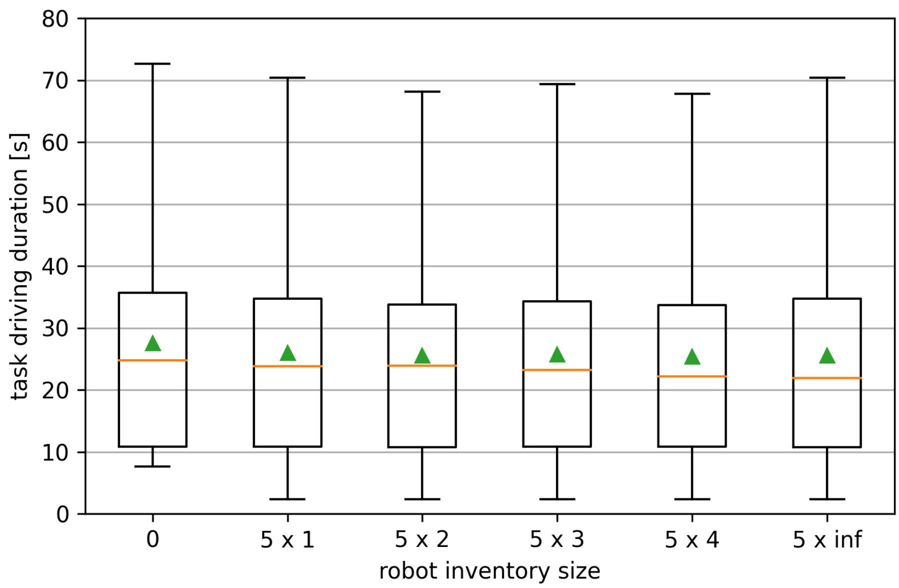

Figure 5.15: Driving durations of intraoperative T1 type tasks are shown as boxplots for scenarios where the 5 most commonly requested articles are held available within the inventory of each robot. The number of units that are stocked per article is varied between 1 and 4. As a baseline for comparison, results for two reference scenarios are given as well: one scenario without robot inventory and one with an infinite number of units per article. Median driving durations are indicated by orange lines, mean driving durations are depicted as green triangles.

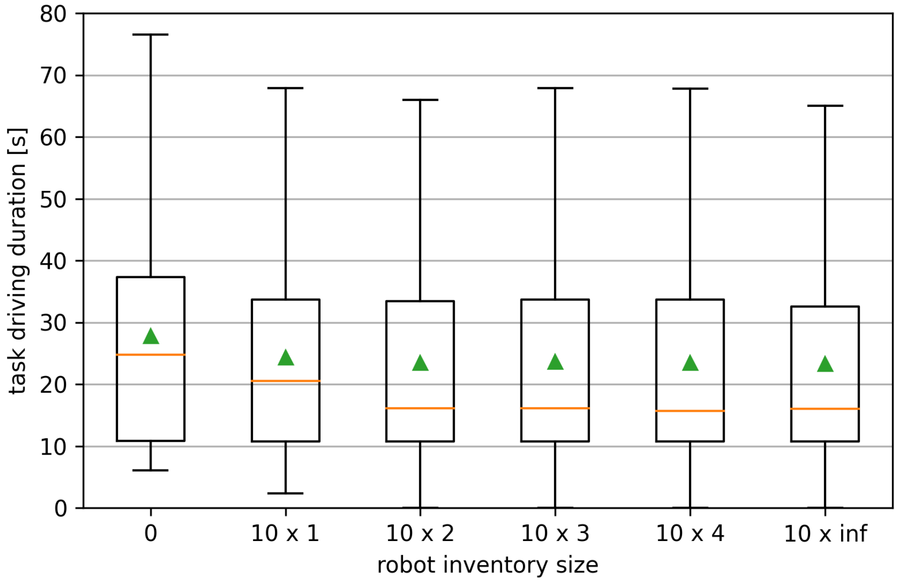

Figure 5.16: Driving durations of intraoperative T1 type tasks are shown as boxplots for scenarios where the 10 most commonly requested articles are held available within the inventory of each robot.

Discussion

The results presented in the previous demonstrate that driving durations can effectively be shortened by the introduction of robot inventories. However, the outcome strongly depends on the dimensioning of the robot inventories. Scenarios with 5 article types show very limited improvement, which should not make a noticeable difference in practice. When stocking 10 article types, the performance improves more obviously, especially when considering the median driving duration. The effects of these improvements would certainly be noticeable in practice. Overall, however, the impact of the robot inventory approach is smaller compared to the quick storage approach presented in section 5.7. Still, it can be seen as a fall-back measure to complement quick storages. In case of false predictions, the correct article may be available in the robot inventory.

As mentioned before, robot inventories stocking up to 20 items are assumed to be realistically achievable within MSR designs similar to the AURORA robot. Such a robot would allow for implementing all studied inventory sizes, except 10 x 3, 10 x 4, 20 x 2. Interestingly, however, these larger sizes do not considerably improve fleet performance according to the presented simulation results. Overall, an inventory of size 10 x 2 can be recommended as a "sweet spot" between space requirements and fleet performance gains.

5.9 Final Approach

In the previous, several key aspects of MSR fleets within the OR wing were studied. This includes fleet dimensioning, fleet composition, battery life, robot velocity as well as approaches for task prioritization, multi-task mission planning, quick storages and robot inventories. Up to this point, these aspects were investigated separately, to be able to better quantify the

Table 5.8: Numeric results for driving durations of T1 type tasks are given. For each simulated scenario, the mean and median driving durations are provided, as well as the relative percentual decrease when comparing both measures to the reference scenario (inventory size of 0). Furthermore, the mean absolute deviation (MAD) is given as a measure for the variance of each distribution. The p-value provided in the last column is the result of the Mann-Whitney U test, when comparing each scenario to its predecessor.

inventory size	mean [s]	mean decrease [%]	median [s]	median decrease [%]	MAD [s]	p
0	27.56	-	24.79	-	14.48	-
5 x 1	25.93	5.89	23.78	4.05	15.00	< 0.001
5 x 2	25.54	7.31	23.93	3.45	14.87	0.262
5 x 3	25.73	6.62	23.18	6.46	15.12	0.382
5 x 4	25.36	7.96	22.17	10.57	14.79	0.362
5 x inf	25.53	7.36	21.93	11.51	15.18	0.439
10 x 1	24.34	12.46	20.54	17.12	14.65	< 0.001
10 x 2	23.42	15.76	16.07	35.16	14.63	0.024
10 x 3	23.65	14.95	16.07	35.16	14.91	0.453
10 x 4	23.44	15.68	15.69	36.71	14.96	0.350
10 x inf	23.30	16.21	16.05	35.23	14.92	0.271
20 x 1	23.04	16.40	18.18	26.66	15.08	< 0.001
20 x 2	22.53	18.26	15.62	36.99	15.67	0.033
20 x inf	22.43	18.61	15.62	36.99	15.69	0.376

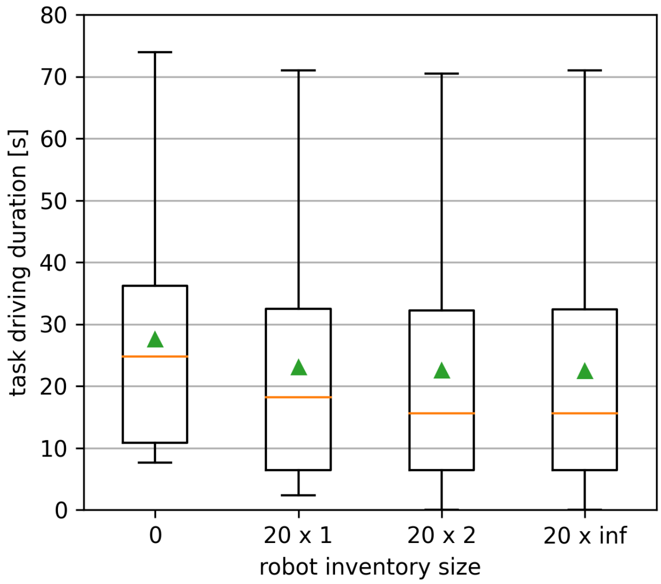

Figure 5.17: Driving durations of intraoperative T1 type tasks are shown as boxplots for scenarios where the 20 most commonly requested articles are held available within the inventory of each robot.

individual benefits provided by each one of them. The simulation study presented in the following takes a different approach and investigates the impact of integrating all concepts and insights gained so far into a final solution. This is further used to compare the outcome to the status quo, i.e. the performance levels reached by human OR teams.

Study Design

For the sake of this study, three simulation scenarios were compared that are described by Table 5.9. Firstly, a human-only reference scenario (ref-human) was simulated by strictly replaying recorded intervention workflows (see section 4.2). Secondly, a scenario was simulated that reflects all insights gained in the course of this doctoral thesis. To that end, the fleet size was chosen according to the optimal value for the given number of operating rooms (see section 5.1). The fleet was composed of all-rounder robots, since a superior performance compared to specialized robots was shown in section 5.2. A zone-based approach was chosen for governing robot speeds, according to the insights gained in section 5.3. Robot batteries were dimensioned in line with the findings presented in section 5.4. Task prioritization was achieved using the ANTS-OR strategy evaluated in section 5.5. The benefits of a multi-task scheduling approach were leveraged, as motivated by section 5.6. Quick storages and robot inventories were dimensioned according to the findings of sections 5.7 and 5.8.

Lastly, a robotic reference scenario (ref-robotic) was simulated which represents the "naive" approach, without applying the findings of this doctoral thesis. This aims to underline the importance of the developed concepts for practical application and for the general feasibility of MSR fleets within the OR wing. While the number of ORs, the fleet size and the fleet composition were kept identical to achieve a fair comparison, no zone-based approach was used for regulating robot speeds, tasks were prioritized using a FIFO approach and concepts like multitask mission planning, quick storages and robot inventories were not utilized.

Table 5.9: The characteristics of the simulated scenarios are summarized. For each approach, the most important simulation parameters are given.

ID	ref-human	own work	ref-robotic
Number of ORs	6	6	6
Fleet Size	0	6	6
Fleet Composition	-	all-rounders	all-rounders
Robot Speed	-	$0.3m/s$ (ORs) / $1.2m/s$ (else)	$0.3m/s$
Battery Simulation	-	yes	yes
Task Prioritization	-	ANTS-OR	FIFO
Mission Length	-	5	1
Article Prediction	-	yes (accuracy: 0.9)	no
Quick Storages	-	yes (threshold: 2)	no
Robot Inventory	-	yes (size: 10 x 2)	no

Results

Table 5.10 and Figure 5.18a show results for the *total task duration*, which denotes the complete duration required for executing a given task, from the time the request was placed by the surgical team to the successful completion of the task. This is a central performance measure from the perspective of the user, since long waiting durations lead to delays within the surgical workflow. As can be observed from the results, there are vast differences between the studied scenarios. The human reference scenario defines the baseline with a mean duration of 56.05 s and a median of 53.01 s. For the final (robotic) approach of this doctoral thesis, an increase by 77.63 % (mean) and by 33.34 % (median) is observed. For the robotic reference scenario, the total task duration increases by extreme values of 1117.46 % (mean) and 413.86 % (median), when compared to the human reference.

Table 5.10: Mean, median and mean absolute deviation (MAD) of total task durations are shown for the simulated scenarios. Furthermore, the relative performance to the human reference scenario (ref-human) is given (mean and median decrease).

total task duration	mean [s]	mean increase [%]	median [s]	median increase [%]	MAD [s]
ref-human	56.05	-	53.01	-	39.30
own work	99.56	77.63	70.68	33.34	61.42
ref-robotic	682.40	1117.46	272.38	413.86	664.13

Table 5.11: Mean, median and mean absolute deviation (MAD) of total intervention durations are shown for the simulated scenarios. Furthermore, the relative performance to the human reference scenario (ref-human) is given (mean and median decrease).

intervention duration	mean [h]	mean increase [%]	median [h]	median increase [%]	MAD [h]
ref-human	2.80	-	2.62	-	0.59
own work	3.02	7.82	2.75	4.89	0.64
ref-robotic	5.54	97.85	5.21	98.93	1.68

Table 5.11 and Figure 5.18b show results for the *total intervention durations* observed for the simulated scenarios. Herein, not only the duration from first incision to wound closure is included, but also the preparation and clean-up phases, since the robotic fleet is heavily involved in these parts of the process. These phases are associated with a high workload and thus may affect the daily schedule of surgical interventions quite considerably in cases of delay. This is especially important from the perspective of the surgical clinic and the hospital as a whole, since fewer patients can be treated within a given timespan, which decreases revenues. For the recorded set of interventions, the human reference scenario yields a mean intervention duration of 2.80 h and a median duration of 2.62 h. For the same set of interventions, the final approach of this doctoral thesis these values can almost be upheld: the simulation results show a slight extension leading to a 7.82 % increase in the mean and a 4.89 % increase in the median. Again, the robotic reference scenario yielded extreme values, indicating a 97.85 % increase in the mean and a 98.93 % increase in the median, effectively doubling the total intervention durations.

Discussion

While the final combined approach of this doctoral thesis does not fully reach the performance level of the human reference, it demonstrates the feasibility of MSR fleets for the OR wing, when applying the insights and approaches described in the previous. In contrast, the outcome of the "naive" robotic reference scenario is clearly not sufficient for real-world requirements. The extreme values observed both for the total task duration and for the total intervention duration indicate a complete overload of the system, with a robotic fleet that cannot keep up with demand at all. In this configuration, the application of MSRs is clearly far from being feasible and instead would lead to a constant frustration of the surgical teams and a strong decline in the number of patients that can be treated within a given day.

When applying the findings presented in previous sections, however, the robotic performance reaches a level that is in the same order of magnitude as the human reference. In particular, the extension of total intervention durations is certainly negligible and should not lead to an unacceptable decrease in patient cases that can be handled by the surgical clinic. Since the robotic fleet is sufficiently effective to not cause a constant overload of tasks, incoming requests can be completed within a reasonable amount of time. While members of the surgical teams will, statistically speaking, have to wait longer, the durations are within sensible bounds and should not cause unacceptable delays within the surgical workflow.

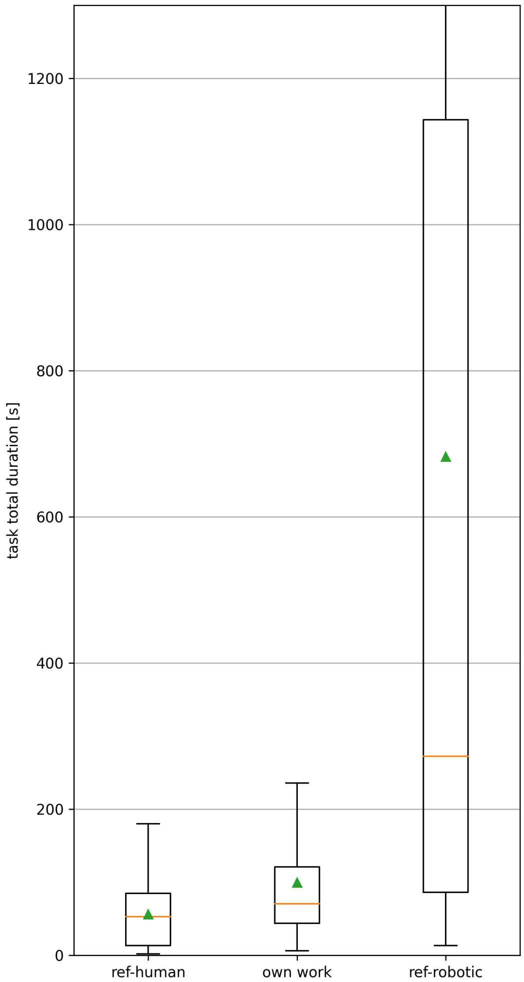

(a) task total durations

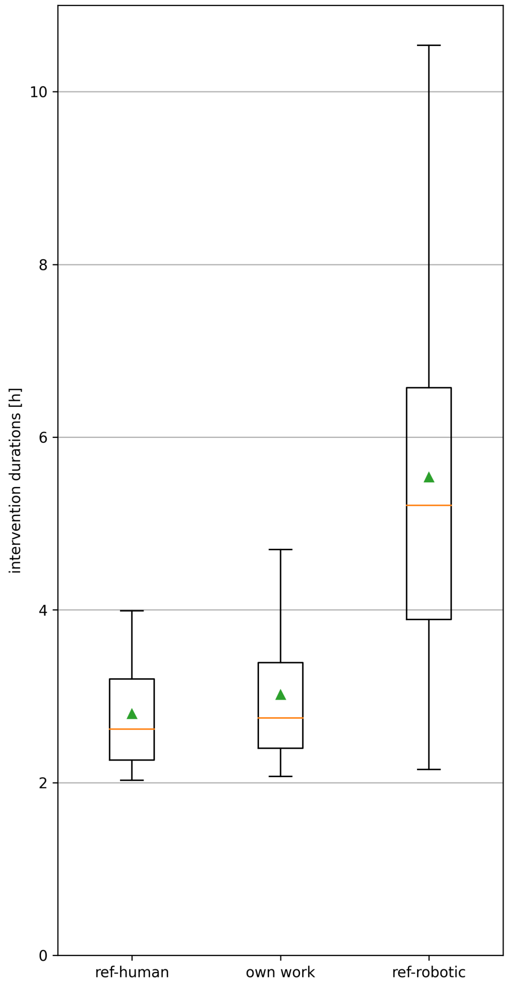

(b) intervention durations

Figure 5.18: Figure part a) shows task total durations observed for the three simulated scenarios. Accordingly, part b) is concerned with intervention durations. Median values are indicated by orange lines, mean values are depicted as green triangles.

6

Discussion

The findings presented in the Results chapter clearly demonstrate that the use of robotic fleets within the OR wing is feasible and beneficial. For the first time, this subject matter was investigated in a structured and scientific manner, providing key insights that will impact the design of mobile service robots and future OR wing environments alike. With reference to the aims defined in chapter 3, substantial insights were gained regarding the unique requirements of the OR wing setting and, based thereon, a fundamental understanding of optimal fleet dimensioning was developed. Several concepts for addressing specific requirements of the OR wing and for further improving fleet performance were presented and evaluated using realistic environmental and procedural simulation models. Building on this extensive methodical and empirical foundation, a final combined solution was proposed and compared to the status quo of a completely human-operated OR wing. The results show that, using the developed methods, key limitations of today's robotic systems can be overcome that would otherwise prevent the successful application of this technology in the surgical domain. By means of adequate domain-specific fleet dimensioning and management, mobile service robotics proves to be highly beneficial for supporting OR wing workflows in a meaningful way. This strongly supports that MSR technology is indeed a powerful tool for alleviating staff shortages within hospitals, and can become one central remedy for upholding clinical operations in the intermediate future.

In the following chapter, the benefits and limitations of the presented work are discussed in detail with respect to the aims and requirements defined in chapter 3. Furthermore, the impact on the interests of different stakeholders is discussed, along with ethical challenges that arise.

6.1 Stakeholders

Multiple stakeholders related to the application of MSR fleets within the OR wing can be identified (see Table 6.1). As primary users of the technology, members of the surgical teams and the OR wing staff are in direct contact with the robotic systems and collaborate with them within the same environment. Interests of these stakeholders include ergonomic working conditions, the avoidance of physical or mental overload, the reduction of less fulfilling tasks, and the reduction of administrative overhead. In particular, staff members expect MSRs to

Table 6.1: The interests of stakeholders that are either directly or indirectly related to the integration of MSRs into the OR wing are summarized.

Stakeholder	Relation	Interests
Clinical Staff	direct	• Ergonomic working conditions • Avoidance of physical and mental overload • Avoidance of less fulfilling tasks • Reduction of administrative overhead
Patient	indirect	• Best possible treatment outcome • Avoidance of physical pain and mental stress • Short and agreeable hospital stay • Personal care, ideally provided by humans
Clinic/Hospital	indirect	• Low costs • Optimal utilization of resources • Staff satisfaction • Positive public image • Long-term plannability
Healthcare System	indirect	• Adequate treatment of all patients • Low treatment costs • Sustainability

reduce their workload in situations of high demand and handle less attractive tasks. They also expect a smooth and non-frustrating collaboration with the robotic assistants.

Other stakeholders take a more indirect perspective, but are also invested in the performance of the robotic fleet: First and foremost, this pertains to the patients being treated by the surgical teams. Their interests include the best possible treatment outcome, a minimum of physical pain or mental stress, and a short and agreeable hospital stay. While, as of now, MSRs do not directly interact with patients, they may be affected by the decision-making of the fleet management and the overall capability of the robotic fleet. If the robotic fleet is temporarily over-challenged, the provision of urgently required materials may be delayed, which in turn may lead to complications and adverse events during the surgical intervention.

The surgical clinic or the hospital as a whole can be identified as a further stakeholder. From this perspective, low operational costs, an optimal utilization of available resources (e.g. operating rooms), long-term retention of personnel and a positive public reputation are favorable. In particular, the hospital management expects robot technology to fill the gaps caused by the staff shortage crisis and operate the available resources to their intended maxima in order to generate sufficient revenues for upholding hospital operations. At the same time, the satisfaction of the remaining personnel is of high interest to avoid frequent staff fluctuations and further shrinking of the teams.

Lastly, the healthcare system and our society as a whole can be identified as stakeholders. From this point of view, ensuring the adequate treatment of all patients, as one cornerstone of our society, is of the highest interest. While this fundamental objective is often taken for granted, it was severely challenged during the Covid-19 pandemic, when hospitals were facing extreme numbers of critically sick patients and many regular surgical interventions had to be postponed in order to cope with the situation. In this context, MSR technology is expected to offer robust assistance around the clock and to flexibly concentrate support where it is currently needed most by means of context-aware proactive fleet management. Further interests of the healthcare system revolve around aspects of sustainability, both to uphold the high quality of patient care in the future and to address new challenges caused by climate change. Here, globally managed robotic fleets are expected to fill in for lacking personnel and improve working conditions for the existing staff, thereby addressing demands of the young generation regarding work-life balance and work-related satisfaction. At the same time, the technology is expected to reduce waste by means of an optimized and need-based material flow.

6.2 Benefits and Limitations

In the following, the benefits and limitations of the presented work is discussed with reference to the aims and requirements defined in chapter 3 and the stakeholder interests described above. Section 6.2.1 covers aspects related to fleet dimensioning, while section 6.2.2 is concerned with fleet management.

6.2.1 Fleet Dimensioning

As a key finding, the results presented in sections 5.1 and 5.3 show that reasonably small fleet sizes suffice for handling the typical workload arising from OR wing operation. In particular, it was demonstrated that a single all-rounder robot per operating room is sufficient for reaching an acceptable fleet performance. This is an important and promising insight regarding the general feasibility of MSRs for the OR wing and is related to the interests of all stakeholders involved. First and foremost, the surgical teams that are working in direct proximity to the robotic fleet and depend on its performance, are not unacceptably disturbed by a large number of robots moving within the confined environment. At the same time, such a small fleet is still capable of handling the imposed workload without unduly delaying the surgical workflow. Long waiting times would hamper the progress of a surgical intervention and likely both frustrate the users and lower their acceptance of MSRs. Knowing the optimal trade-off between fleet size and fleet performance is also highly beneficial from the perspective of the surgical clinic and the hospital, since fewer systems must be acquired, sustained and maintained, which benefits the economic feasibility of this technology. Larger fleet sizes naturally come at higher costs, which could be a serious hurdle for real-world application, even if the effectiveness is clearly evident. For the healthcare system and society as a whole, the feasibility of this technology with regard to required fleet sizes shows its potential for combating current and future nursing shortages.

However, to reach the required performance levels while using small fleet sizes, task execution durations must be kept low. While the most straightforward way to achieve this is by using fast-driving robots, the results presented in sections 5.1 and 5.3 show that at least a driving speed of $v_r = 1.2m/s$ would be required. Such high speeds would certainly be difficult to realize, at least for the complex, dynamic, confined and safety-critical spaces within operating rooms. In the previous, several measures for dealing with this problem were proposed and evaluated. This includes means for decreasing the driving durations, such as the composition of multi-task missions from individual requests (section 5.6), the use of quick storages (section 5.7) and the introduction of robot inventories (section 5.8). As confirmed by simulation, all of these approaches contribute to a significant decrease in task execution durations. This is beneficial from the perspective of the surgical teams, since waiting durations are shortened and a smooth surgical workflow is ensured. Furthermore, shorter task execution durations result in shorter surgery durations, which benefits economic interests of the clinic and allows for the treatment of more patients. The latter is beneficial from the perspective of individual patients, since waiting times for surgery appointments are decreased, and also beneficial from a societal perspective, since the overall capacity of the healthcare system is improved.

As a further measure, the introduction of speed zones was studied, where robot driving speeds are varied across different parts of the OR wing (see section 5.4). By prescribing a considerably slower robot velocity of $v_r = 0.3m/s$ within operating rooms and allowing higher speeds only within the remaining environment, a fleet performance level was observed that is still much closer to the one achieved for fast-only scenarios than for slow-only scenarios. Slow driving speeds within operating rooms are beneficial for all present persons, since their safety level is elevated, the perceived threat of fast-moving robots is decreased and other confounding factors, such as driving noises and vibrations, are improved as well. As a downside of the introduction of speed zones, structural changes to the environment are likely required, such as robot driving lanes on the corridors, similar to those found in industrial manufacturing facilities. For existing facilities, these may be difficult to implement. However, the optimization of architectural aspects may also offer new opportunities for achieving further improvement of MSR fleet performance. This is especially relevant for the design of future hospital facilities, where, for example, dedicated driving passages for mobile robots could be introduced or the layout of the OR wing could be optimized for reducing the length of robotic driving paths. This is strongly connected to the work presented by Amato et al. [Ama+21; Ama+22] and can be considered one prime example of the *surgineering* paradigm [Feu+19].

Further valuable insights have been gained regarding fleet composition. Based on the findings presented in section 5.2, it is beneficial to design robots that can execute a wide range of different tasks. For specialized robots, a significantly worse outcome was observed with regard to global fleet performance, which is why larger fleet sizes are required to achieve a similar performance to all-rounder robots. While all-rounder robots are desirable from a fleet management standpoint, this approach poses challenges to the design of individual robotic systems. Robots with a variety of capabilities tend to be larger, more complex, more costly and less robust. However, it is argued that these technical challenges must be accepted and overcome in order to be compatible with the OR wing setting. With respect to the interests of the involved stakeholders, as described in the previous, it is argued that the positive effects of keeping the fleet size small outweigh the benefits of a specialized robotic design.

According to the findings presented in section 5.4, the capabilities of modern battery systems for mobile robots are generally sufficient for the OR wing scenario. By means of adequate battery dimensioning, forced recharges during the intraoperative phase of surgical interventions can be avoided and robots can sustain their battery power only by intermediate recharging during situations of low demand. Depending on the average power consumption per task execution, which in turn depends on robotic driving speeds and other factors, the required battery capacity varies. While large battery systems are needed to cope with a higher average task power consumption, the requirements are still within technically feasible bounds for the simulated scenarios. By avoiding forced recharges, a better availability of the robotic fleet can be achieved, which in turn leads to a smoother workflow, shorter task execution durations and shorter intervention durations, with all the stakeholder benefits associated with that, as described in the previous.

6.2.2 Fleet Management

As demonstrated in section 5.5, a situation-dependent prioritization of user requests can be achieved by means of the ANTS-OR strategy. Using simulation, it was shown that the resulting response times (idle durations) of tasks can indeed be reliably modulated according to the priorities prescribed by the ANTS-OR score. Thereby, a more meaningful execution order is achieved with respect to the individual importance of requested tasks. This is beneficial from the patient's perspective, since tasks related to patient well-being are executed with preference over less critical tasks. This behavior is also expected by the surgical team, along with the prioritized execution of tasks originating from higher-ranking team members. By realizing a more logical execution order of tasks, the overall workflow is optimized and user acceptance of the robotic fleet should improve as well. At the same time, it was shown that other global fleet performance parameters, such as average intervention duration, are not considerably worsened by the introduction of ANTS-OR prioritization, which means that the economic and societal interests of the other stakeholders are not compromised compared to FIFO prioritization.

However, one challenge arising in the context of ANTS-OR is the automated inference of appropriate command and emergency levels. Such ratings must be available for each task request as a prerequisite for calculating the score. There are several approaches to how this can be achieved for command levels. For tasks originating from digital systems, the command level of *AI-0* can be assigned automatically by the system itself. For tasks requested by human members of the team, the speaker (and thereby the appropriate command level) could be automatically identified by means of voice analysis. While this would be the most intuitive solution, a robust speaker recognition software is necessary, which may require the surgeon to wear a headset. Another possibility is to use a graphical user interface to place task requests, for which the user must be logged in to the system, similar to table management software used in restaurants. The automated inference of appropriate emergency levels is even more challenging. Here, the condition of the patient is the decisive endpoint, which is difficult to quantify objectively, especially without the input of a human expert. The following approaches are proposed: Based on available input parameters – such as vital signs, medical history and preoperative findings – an algorithm rates the criticality of the patient's condition. If the algorithm yields that the patient is in a life-threatening condition, the highest emergency level is assigned to all tasks requested until this situation ceases. If the algorithm yields that the patient's condition may have adverse health consequences without likely being fatal, the intermediate emergency level is assigned. Clearly, such an autonomous decision-making algorithm is challenging to implement and may introduce ethical predicaments, as is further discussed in section 6.3. On the other hand, it would be able to globally consider all available information regarding the state of different patients within a given OR wing. In principle, this puts the algorithm in a much better position for reaching a fair compromise than a human decision-maker who is only able to consider "local" information and may be biased in that respect. As an alternative approach, the surgical intervention can be analyzed using real-time workflow recognition algorithms, such as [Cze+21]. Deviations within the standard workflow may indicate that something went wrong, such as the extensive use of the suction/irrigation device, which may be related to the presence of a severe bleeding. A central challenge of this approach is that not all deviations from the standard procedure are critical and the algorithm must be able to reliably discriminate between those that are and those that are not. Sometimes, unexpected additional tasks must be carried out during a surgical intervention, such as the removal of tissue adhesions. These are usually noncritical and should not be associated with an elevated emergency level. Further hints regarding the urgency of a given situation may be gained from observing the behavior and the emotional state of the surgical team. Very hectic situations – that may be associated with short imperative voice commands, elevated heart rates, elevated breathing rates or increased speaking volume – are strong indicators for the occurrence of adverse events impacting the patient's condition. As a fallback solution, the surgical team could simply be asked to rate the emergency level when making the task request. For instance, this could be achieved by including such a rating within the voice command, when making the request: "*Fetch object X with highest priority*" or "*Fetch object X with elevated priority*". In cases where a rating is not given, the lowest emergency level is assumed. Similarly, a desired emergency level could be assigned when using a graphical user interface. Assuming that the input of the user is qualified and the emergency level is not exploited for unjustifiably raising the priority of a non-critical task, this approach arguably represents the most reliable emergency level rating, since it conveys the intent and expertise of the surgical team. As a drawback, this approach generates additional overhead for the surgical team when placing task requests. Also, as discussed in the previous, the rating may be subjective and would mainly consider "local" information.

As a further major contribution, the VRP-OR problem was presented in section 4.4.3. This variant of the classical VRP is unique to the OR wing scenario and has not yet been described in scholarly literature, to the best knowledge of the author. The VRP-OR can be understood as a combination of different known VRP variants with further additions to address the characteristics of the OR wing, such as alternative pickup locations and prioritized tasks. In the form of the VRP-OR, an essential framework for the management of MSR fleets within the OR wing is now available and will serve as a foundation for future work. The results in section 5.6 demonstrate that the proposed greedy solution algorithm performs well and considerably reduces driving durations during multi-task robot missions. At the same time, the required computation effort, and thus the processing time for performing a full scheduling cycle, is well within the acceptable limit of 800 ms prescribed by human reaction time. As shown in section 5.6, the computational effort scales linearly when increasing the size of the problem instance with respect to mission length, number of task requests or fleet size.

A drawback of the current mission planning approach is the fixed mission length. In some situations, it may be advantageous to split longer missions into multiple shorter ones in order to ensure the timely execution of high-priority tasks. Consider the following exemplary mission of length 8 consisting of two tasks with the highest ANTS-OR emergency level of E-2, followed by six further tasks with a low emergency level of E-0:

Mission 1							
E-2	E-2	E-0	E-0	E-0	E-0	E-0	E-0

Here, the high-priority emergency tasks were assigned to the first spots of the mission to minimize the idle duration for these urgent requests. However, following the proposed algorithm, the other spots were filled with low-priority tasks (assuming no further emergency tasks were requested or they were already assigned to other robots). This behavior is disadvantageous from the standpoint of the surgical team dealing with the emergency, since the mission is elongated due to the combined execution with low-priority tasks. Thereby, the completion of the emergency tasks is avoidably delayed. It is argued that the scheduling algorithm should instead split the original mission into one critical and one non-critical mission:

Mission 1a	
E-2	E-2

Mission 1b					
E-0	E-0	E-0	E-0	E-0	E-0

By means of such a dynamic decomposition of long missions, the delayed execution of emergency tasks could be avoided. In more complex cases, however, where many different priority levels are present within a given mission, the decomposition may not be as straightforward as in the above example, since task priority must be adequately balanced with overall execution time. Thus, the development of suitable strategies for context-dependently decomposing missions remains to be addressed in future work.

In sections 5.7 and 5.8, two proactive material flow concepts were presented that aim at further improving the efficiency of surgical workflows. Firstly, robots can prepare for upcoming interventions by holding the most commonly requested articles readily available using built-in inventories. Secondly, additional quick storages can be introduced into the environment and proactively stocked with articles that are likely needed within the immediate future. According to the presented simulation results, both approaches lead to a considerable reduction of robot driving times, which in turn reduces waiting times of the surgical teams and leads to shorter overall durations of the interventions, with all the associated stakeholder benefits, as described in the previous. The two presented approaches are model examples of how the introduction of MSRs provides opportunities for re-thinking current material flow concepts within the OR wing. By means of providing packaged sterile goods on an "as needed" basis – instead of tentatively preparing and opening them preoperatively – the material expenditure can be reduced, which benefits both hospital and society, since overall treatment costs are decreased and sustainability is improved.

6.3 Ethics

In the following, ethical predicaments arising from task prioritization in the context of the OR wing are identified (section 6.3.1) and first solution strategies are presented (section 6.3.2). This is complemented by an ethical discussion of whether the idea of MSR fleets for the OR wing should be embraced or rejected from a consequentialist perspective (section 6.3.3).

6.3.1 Ethical Predicaments

As described in section 4.4, one core part of fleet management is the allocation of tasks to MSR resources. Therein, tasks may have different priorities and robotic systems may be scarce, which immediately poses the question of how one can make the "best" use of what is available. For similar multi-robot task allocation problems found in other domains (such as industrial manufacturing and logistics), the answer to this is often simply "cost efficiency". In the healthcare context, however, other endpoints, most importantly patient well-being and outcome, are of higher priority. Thus, when distributing tasks to robotic resources, the impact of these decisions on the patient must be considered. This is unproblematic for situations concerned with tasks that are non-critical regarding patient health or when enough robotic resources are available to sufficiently handle all requested critical tasks. Clearly, robotic fleets should be adequately dimensioned to be able to cope with the intended workload resulting from normal (i.e., common and planned) clinical operation, in accordance with the results presented in section 5.1. However, it can never be fully guaranteed that demand can always be met, at which point problems start to arise. Such situations may be due to unexpected events resulting in very high patient numbers, such as mass casualty incidents, and may be facilitated by external factors, such as temporary understaffing of the clinical personnel. In such situations, robotic resources may not be able to handle all critical tasks in time and the fleet manager must decide between the needs of multiple patients being simultaneously treated in different operating rooms. While these situations are rare and represent extreme cases, they will arise eventually and thus there is a responsibility to develop methods to resolve them ethically, as a prerequisite for integrating MSR fleets into real-world OR wings. Clearly, such considerations may become arbitrarily complex, especially when pondering different patient scenarios and conditions.

The following simple example illustrates the ethical predicaments that may arise in some exceptional situations: Two surgical teams in two separate ORs are both facing a critical situation where certain materials (e.g., banked blood) are urgently required. Task requests for fetching the desired materials are requested by both teams simultaneously and thus reach the fleet manager at the same time. In case robotic resources are currently overburdened, the fleet manager may have to decide which OR, and thus which patient, should be served with priority. While this is a rather naive example and may seem hypothetical, it illustrates which kind of dilemmas may arise when an overload of urgent tasks has to be ethically allocated to limited robotic resources – thereby directly impacting the well-being of different patients. Resolving such conflicts may be challenging even for humans, and more so for automated systems. However, temporary overloads of the available resources due to unexpected events can never be completely ruled out, which is why methods must be developed for handling them ethically.

6.3.2 Resolution Strategies

With reference to the principles used in clinical recommendations for the admission of Covid-19 patients to intensive-care units [Deu20], it is argued that aspects such as age or pre-existing diseases should not be of influence. Instead, only "ad hoc" parameters should be considered, such as the current survival probability or the amount of health damage caused by not executing a task on time. However, estimating these parameters is extremely difficult and the available information is limited. In the digital operating room of the (near) future, it can realistically be expected to have instant digital access to the following information:

- Patient record/history (age, pre-existing diseases, reports etc.)
- Type of surgery being performed
- Vital signs
- Additional sensors within the OR (e.g. cameras)

This information can be used to rate the urgency of a given situation using a dedicated scale, where each level is delimited by distinct criteria. This approach is inspired by triaging procedures used for the initial assessment of crash victims (e.g., START Triage Algorithm [US 24]), for emergency admission to the hospital (e.g., Emergency Severity Index [GDK09]) or for admission to the intensive care unit [Deu20]. Such procedures usually rely on clearly defined measures (such as clinical frailty scores or vital signs) and a classification scheme to prioritize patients with regard to these measures. However, it is important to stress that in all these procedures the final decision is always made by a qualified human expert, not an autonomous technical system. This is different when considering control algorithms for self-driving cars. Here, the algorithm may also run into situations where it has to decide between two potentially disastrous options (e.g. crashing into car A or car B), based on limited sensor data and without human assistance. In this context, predicaments are resolved by considering all courses of action and calculating risk magnitudes for each one. Answering the questions of how available information can be used for calculating such risk magnitudes in the OR wing context, and whether this information is sufficient for ethical decision-making, will require extensive further considerations in the future, framed by in-depth interdisciplinary discussions.

As a means for effectively yielding the responsibility back to qualified humans, clinicians could be asked to provide an urgency rating when issuing a task request – e.g. "Fetch blood bags with urgency 5". Using this approach, a coarse prioritization of tasks can be achieved, which is presumably sufficient to resolve many real-life situations. However, such ratings are bound to be partly subjective and most likely biased with respect to the specific situation the clinician is facing. Furthermore, special cases remain where two tasks with identical urgency ratings have been issued at the same time, while not enough robotic resources are currently available to execute both in parallel.

Since critical situations leading to ethical dilemmas are expected to be rather exceptional during regular clinical operation, it may also be sensible to take a two-staged approach: In this context, the automated system would be designed to merely deal with "regular" situations, where task scheduling is not directly impacting patient health. This should be sufficient for dealing with most situations occurring during everyday clinical operation. In the case of a critical "non-regular" situation, the system would then yield responsibility to a human decision-maker. While this approach circumvents the necessity of automatically resolving ethical dilemmas, other challenges arise: Firstly, the system must be able to adequately assess, what qualifies as a “non-regular” situation and, secondly, ways (i.e. interfaces and routines) need to be found to involve a human decision-maker into the task prioritization process.

A further approach for circumventing the automated resolution of ethical dilemmas is to always assign tasks with emergency level E-2 to humans instead of robots. Thereby, the responsibility for further prioritizing these tasks with respect to each other remains with human decision-makers. As a drawback, the robotic fleet now relies on humans as a backup for emergency situations, which requires an adequate number of persons to be available at all times.

6.3.3 Ethical Discussion

In the following, a consequentialist approach is used to discuss whether an ethical obligation to embrace or reject MSR fleets for the OR wing can be inferred.

"The application of MSR fleets within the OR wing is unethical"

As of now, critical decisions within the OR wing are made by well-trained human experts who can rely on extensive background knowledge, both theoretical and empirical, as well as a conscience shaped by education and society. Both of these aspects cannot be instilled into a technical system, at least to its fullest extent, using the technology available to us today or in the foreseeable future.

At the same time, it is implicitly accepted that human beings sometimes make mistakes, even when acting according to best knowledge and conscience. The necessary societal tools for handling these cases, such as insurance and jurisdiction, are established and available to us. The question of responsibility can usually be resolved to one or several of the involved persons. There is a defined, albeit often cumbersome and time-consuming, process to work out appropriate legal actions and compensations. Regarding autonomous technical systems, on the other hand, matters are not yet as well-defined. Similar to the issues arising in the context of self-driving cars, it is unclear, who can be held responsible for consequences of autonomous decision-making, which can be seen as reasons to reject the idea for now.

"The application of MSR fleets within the OR wing is ethical"

The described ethical predicaments only arise in rare and extreme cases, when facing workloads for which the system as a whole has never been dimensioned in the first place. It can be argued that especially in these situations robotic assistance is most welcome and required, since there is no human worker available to take the robot’s spot. Therefore, it can be inferred that we have an ethical obligation to make use of this tool to benefit the overall outcome, even if its decisions are not ethically impeccable. This, of course, assumes that the robot-assisted case does not lead to worse results than a case without additional robotic (nor human) support.

Furthermore, having a global perspective on OR wing processes – enabled by networking OR resources and by introducing an orchestration algorithm – allows for optimizing the entire system as a whole, leading to a more economic, ethically fairer allocation of resources to points of need. This will allow for treating more patients and also improve the quality of treatment since urgent needs can be addressed with priority and with higher workforce. Based on this, it can be argued that there is an ethical responsibility to make these benefits available to the patient. This is in contrast to the traditional approach, where a single OR is mostly decoupled from the decision-making of other ORs.

Lastly, we have a responsibility to uphold the capacity of our healthcare systems as well as today's excellent treatment quality, which are both severely threatened by long-term staff shortages, rising patient numbers and over-aging. In cases where we simply do not have the personnel (or other resources) to treat all patients, we are obliged to make use of any support that is available to us, including mobile service robots, even if their performance may not yet fully live up to human healthcare workers.

Recommendations

Since the necessary workforce for upholding our current standard of healthcare is lacking, the use of service robotics can hardly be avoided, despite the risks and limitations involved. To make the best use of this technology, global orchestration strategies are required, which may involve deciding ethical predicaments in an automated fashion. In this regard, a decision-making performance that is inferior compared to human performance can be accepted from a consequentialist perspective, since the alternative option of having no robotic support threatens the functioning of our healthcare systems. From a societal standpoint, this clearly must be strongly avoided. However, new ethical and legal frameworks must be developed to govern the future advancement of MSR technology for the OR wing and address the questions of responsibility and liability.

7 Conclusion

Due to pressing problems such as increasing patient numbers, over-aging and shortage of qualified workers, the sustainability of today's healthcare systems is severely challenged. As a complement to urgently needed societal and political reforms, available technologies must be exploited in order to uphold and further improve the current quality and capacity of patient care services. In this context, the application of assistive robotics is highly promising and holds great potential for supporting clinical workflows and for relieving overburdened personnel.

The operating room wing is a clinical unit that is particularly affected by the lack of qualified workers, which increasingly leads to the postponement or cancellation of surgical interventions. As demonstrated by the AURORA research project, mobile service robots are able to support surgical teams by executing frequently required tasks suitable for robotic automation. However, the challenging domain-specific requirements and environmental characteristics limit the effectiveness of individually acting robots. In the framework of this doctoral thesis, the potential of combining multiple robotic systems into globally orchestrated fleets was studied for the first time. Central aims were the development of guidelines and concepts for the dimensioning and composition of MSR fleets as well as the development of fleet management techniques suitable for the surgical domain.

In order to investigate diverse fleet arrangements and validate proposed concepts, a simulation-based methodology was employed, facilitated by the implementation of a novel simulation framework referred to as FleetOR which represents a valuable contribution in its own right. FleetOR allows for the simulation of MSR fleets within OR wing environments while exposing them to authentic workflows and workloads. MSR fleets can be freely composed in order to study the strengths and weaknesses of different fleet arrangements and sizes. As a prerequisite for the conducted simulations, a model OR wing was implemented, mirroring the facilities of a real German university hospital. Through observational studies, the workflow of surgical procedures was documented, and specific types of tasks suitable for mobile robotic execution were identified.

By means of the FleetOR simulation framework, multiple investigations were carried out, exploring crucial elements of MSR fleet management. In a first step, the appropriate dimensioning of robotic fleets was studied. This investigation aimed at providing insights regarding the optimal number of robots per operating room – a critical factor for efficiently managing real-world workloads. Results of the study indicate that the ideal fleet size is closely tied to the speed at which the robots operate. For high driving speeds ($1.2m/s$), a single robot per OR is sufficient, but slower driving speeds ($0.3m/s$) necessitate a considerably higher number of 3-6 robots per OR. Given the challenges of achieving consistently high driving speeds in the intricate, dynamic, and safety-critical OR wing environment, strategies were explored to overcome this hurdle. It was shown that by implementing zones of different driving speeds, a viable balance between speed and overall performance can be reached. As a further central finding, employing versatile, allrounder robots over specialized robots was found to result in a vastly superior fleet performance.

As a further key contribution, domain-specific fleet management concepts were developed that reflect the specific requirements of the OR wing. To that end, the novel ANTS-OR scoring strategy was introduced, as a means for assigning task priorities based on the current clinical context. This prioritization method takes into account domain-specific knowledge, such as patient condition and clinical command hierarchy. The effectiveness of the ANTS-OR task prioritization strategy was assessed in comparison with a first-in-first-out methodology. The findings demonstrate a distinct advantage in handling emergency situations: With the implementation of ANTS-OR, crucial tasks related to patient well-being are executed significantly earlier than less urgent tasks.

The optimal composition of individual task requests into comprehensive multi-task missions is governed by the *Vehicle Routing Problem for the Operating Room Wing* (VRP-OR). As a key contribution of this doctoral thesis, the VRP-OR was motivated and formally described for the first time. Given the NP-hardness of this problem, a heuristic algorithm was proposed to solve instances within an acceptable processing time, allowing for a quick adaptation of robot mission schedules to changing circumstances – a crucial requirement in the surgical domain. When assessing the performance of this algorithm, significant decreases in task driving times were observed for multi-task missions as compared to planning task executions in an isolated fashion (48.7 % decrease for a mission length of 4). Compared to a computationally taxing calculation of the optimal solution, only a small average increase in driving duration (6.8 %) must be accepted when using the heuristic solver. At the same time, the processing times observed for the heuristic solver were exceptionally brief, with an average of less than 38.9 ms. To further improve the performance of robotic assistance, the merits of a proactive behavior of MSR fleets were studied. This involved anticipating the intraoperative use of surgical materials to prepare likely future requests in advance. In this context, logistical concepts such as quick storages and robot inventories were studied. Notably, the implementation of quick storages has demonstrated significant potential in reducing task execution times, even for lower prediction accuracies.

Finally, all prior findings and concepts were combined into a comprehensive approach, which was benchmarked against both a human-only scenario and a "naive" robotic reference scenario. An acceptable performance of the robotic fleet was only achieved by incorporating the results of this doctoral thesis. In contrast, the naive reference scenario proved inadequate for real-world requirements, exhibiting an inferior outcome by one order of magnitude. These results strongly demonstrate that integrating MSR fleets into surgical OR wing environments is indeed viable and beneficial. By consequently applying the findings of this doctoral thesis, the integration of MSR fleets into the OR wing becomes feasible for the first time. Limitations of today's robotic systems, such as driving speed, manipulation speed, or battery capacity, are compensated by employing novel techniques for fleet dimensioning and fleet management. The presented simulation results show that realistic workloads can be sufficiently handled, such that the current patient throughput is sustainable. By addressing essential domain-specific requirements, such as the prioritization of competing tasks and the need for short response times, the applicability of the developed concepts to real-world environments and workflows is ensured.

When considering the perspectives of key stakeholders such as clinical personnel, patients, hospitals and the healthcare system, it can be concluded that the adequate implementation of MSR technology yields advantages for all parties involved. Staff experience benefits from the compensation of personnel shortages and from a reduction of undesirable or non-ergonomic tasks, thereby fostering a healthier work environment and mitigating stress and burnout. Patients benefit from caregivers having more time for direct care and from an enhanced capacity of surgical services, resulting in fewer delays and cancellations of planned surgeries. Hospitals gain advantages by offering more appealing working conditions, thereby improving long-term staff retention, and by increasing revenue through the treatment of more patients. Overall, the healthcare system benefits from an enhanced resilience of patient care services, even during periods of high demand.

From a more general perspective, the concepts presented within this doctoral thesis introduce a fundamental paradigm change in the organization and management of modern OR wings. As of yet, processes are mostly managed on the level of single operating rooms, each representing an individual and largely encapsulated unit with dedicated human and physical resources allocated to it. While high-level planning and guidance are typically provided by a central OR wing control center, there is no central instance for monitoring and controlling the global workflow and material flow on a fine-granular level. While the drawbacks of this federated approach may not necessarily be obvious from the perspective of single ORs, available resources are often not fully utilized, which becomes apparent when taking a global perspective. By introducing a centrally orchestrated MSR fleet into the OR wing, a fundamental transition from an OR-individual to a global management of OR wing processes can be achieved. This offers the chance to globally optimize the workflows, assign available resources to sites where they are currently needed most and provide materials in a more sustainable need-based manner. Clearly, such a global optimization of OR wing processes benefits all stakeholders and is certainly also applicable in a more generic context beyond MSR fleet management.

8

Future Work

In the previous, the benefits of MSR fleets for the OR were demonstrated and concepts for a domain-specific use of this technology were presented and validated. Based on ethical considerations, an obligation to employ MSR fleets as one central remedy for addressing pressing problems revolving around staff shortages was identified, as a means for sustaining the high standard of today's surgical patient care. However, further work needs to be conducted to facilitate the successful translation of this promising technology into real-world hospitals. These remaining limitations and challenges will be addressed by future research and development efforts, for which several focus areas are outlined in the following.

Fleet Management

The VRP-OR, which was introduced in the previous, provides a framework for mobile robotic fleet management within the OR wing. In the future, its formulation will be extended to include further aspects, such as the shared execution of tasks, where multiple robots collaborate to achieve an overarching goal. This allows for realizing more complex task workflows, which may require robotic capabilities that are too diverse to fit into a single system. This notion, which is referred to as multi-robot tasks (MR) according to the taxonomy introduced by Gerkey and Mataric [GM04], has already been studied in other domains [Cha+23] and can potentially be adapted to the OR wing scenario. Similarly, managing the collaboration between human workers and robots can be explicitly included in the scheduling, with some parts of a task workflow being executed by humans and others by robots. This leads to the problem of optimal human-robot task distribution based on individual strengths and weaknesses, which, again, has been studied in other contexts (e.g., [BKW20]) but must be adapted to reflect the requirements of the OR wing.

In the context of advancing fleet management strategies for the OR wing, the ongoing use of simulation will facilitate the optimization of workflows and help to develop and evaluate new concepts. The novel simulation environment FleetOR proved to be a highly valuable means for studying and comparing different approaches and MSR fleet configurations in the context of this doctoral thesis. In the future, this framework will be further extended and generalized, as a natural complement to the advancement of fleet management concepts for the OR wing.

Workflow Anticipation

According to the results presented in the previous, the prediction of future task requests is an effective strategy to improve overall fleet performance and compensate for limitations of individual robotic systems regarding execution speed. Such a prediction can be achieved based on structured domain knowledge, sensor-based inference or a combination of both. By processing available data originating from sensors, devices, and clinical information systems and augmenting this data with structured knowledge originating from models and onotologies, algorithms can infer the current situation within the OR (wing) and predict situations that are likely to occur in the future. These current and future situations may be associated with certain needs that must be served and tasks that must be executed. If possible, these tasks can then be prepared or fully executed by the robotic fleet ahead of time, instead of waiting for a specific request by the surgical team. Situational information can further be used to optimize robot traveling routes and idle positions in order to respond faster to arising requests. Thus, a mere reactive behavior of the robotic fleet can be advanced towards an anticipative behavior, which is the prerequisite for achieving reaction times compatible with real-world requirements.

In scientific literature, approaches for tracking the surgical workflow on a coarse-granular level (phase recognition) have been presented [Cze+21] as well as means for recognizing low-level actions of the surgeon [Pet+18]. Both tasks are commonly achieved based on processing the video stream of the laparoscopic camera unit but can be augmented by other data sources. Further work exists regarding the prediction of the remaining surgery duration [Twi+19], which can be used to optimize the planning horizon of mobile robotic scheduling. Not merely tracking but *predicting* individual steps of the workflow or, more directly, specific needs of the surgical team, may prove even more useful. One existing example is an ontology-based approach for predicting workflows in neurosurgery, which was proposed in [Neu+22]. Furthermore, in the framework of the ongoing research project SASHA-OR [Wag+22; Wag+23a; Wag+23b], a multi-modal framework for intraoperative inference of contextual information is currently in development. As of now, this is used to predict the intraoperative use of laparoscopic instruments, which enables the robotic automation of handing over instruments to the surgeon. While this use case takes place completely within the sterile zone, and thus is not related to non-sterile mobile service robotics, the framework can be extended in the future to predict further instruments and materials required during surgical interventions, including single-use materials stored outside the sterile zone. By predicting the most likely materials that are needed next in addition to the time horizon in which they will be needed, the fleet management system can schedule robots to fetch the materials ahead of time and place them in their inventory or already transport them to the OR.

Capabilities of Individual Robots

Future research efforts must further be concerned with improving the capabilities of individual robotic systems. As was confirmed by the simulation results presented in the previous, the design of versatile robots with a range of different capabilities is highly beneficial from the perspective of fleet management and can even be identified as a requirement for achieving acceptable performance levels with smaller fleets. So far, existing systems, such as the AURORA robot, are rather limited with respect to the types of tasks they are capable of executing. As of now, AURORA focuses on tasks that, according to observational studies conducted in the framework of the project, are requested most frequently during the preparation or intraoperative phase of surgical interventions: The provision of sterilely packaged surgical materials and the adjustment of medical devices. Furthermore, AURORA is limited with respect to the types of materials it can handle. The current prototype focuses on suturing material and surgical gloves, which represent the most frequently required materials. However, depending on the type of surgery being performed, many more materials may be relevant. Moreover, material provision and device adjustment may account for a large portion of requests being made by the surgical team, but they represent only a small subset of the duties of human circulating nurses. While some of these task types are better addressed by other technological means (digitalization), others are promising for mobile robotic automation and should be considered by future development efforts. Most notably, this may refer to the transportation of large objects (devices, instrument sieves, patient beds), assistance of the surgical team during the sterile dressing procedure, counting and disposing of used materials, transporting and tube-mailing of biopsy samples, cleaning of surfaces and disposal of waste.

Hybrid Human-Robot Work Environment

Beyond improving the performance and safety of individual robotic systems, adequate peripheral conditions must be created for them to operate well. Surgical environments must be transformed towards being better suited for robotic requirements while remaining fully compatible with the needs of human workers and patients. This includes the redesign of storage facilities, surgical layouts, and architectural aspects. As of now, storage solutions within OR wing facilities are tailored toward use by human workers and, thus, are not necessarily compatible with robotic requirements. For instance, materials may be stored in loose heaps within drawers or behind doors, which is challenging for robotic interaction. In order to further improve robotic execution times and robustness, storage cabinets must be redesigned to be compatible with the limitations of typical robotic systems with respect to degrees of freedom, workspace, and perception. Further considerations must be made regarding the room layouts used during surgical interventions since mobile robotic systems require space to move through the environment and need to access points of interest during task execution. In this context, clinical guidelines must be revised to include aspects such as the clearance of driving passages as well as the management of cables and tubes within the OR. Such considerations must also be extended to peripheral rooms and corridors within the OR wing, where the movement of robotic systems must be coordinated with the movement of human workers and the storage of large equipment. The introduction of dedicated driving lanes for mobile robots or fully encapsulated corridors may be investigated in this context.

Since human workers will remain essential for most workflows taking place within the OR wing, the facilities must be transformed into a shared work environment in which both human and robotic agents find suitable preconditions for living up to their fullest potential. Harmonizing the requirements of both sides to create a hybrid human-robot work environment is an intriguing architectural, mechatronic, and logistical problem to be solved in future work. Here, the *surgineering* paradigm [Feu+19] will be a central instrument to develop meaningful solutions and rapidly translate them to clinical practice. First work in this context has been presented by Amato et al. [Ama+21; Ama+22], who aim to develop architectural hospital layouts that are human-centric yet compatible with existing and upcoming technological solutions.

List of Figures

List of Tables

Bibliography

[AIS20] Abrar, M. M., Islam, R., and Shanto, M. A. H. “An Autonomous Delivery Robot to Prevent the Spread of Coronavirus in Product Delivery System”. In: *2020 11th IEEE Annual Ubiquitous Computing, Electronics & Mobile Communication Conference (UEMCON)*. IEEE, 2020, pp. 0461–0466. ISBN: 978-1-7281-9656-5. DOI: 10.1109/UEMCON51285.2020.9298108.

[AMN15] Acosta Calderon, C. A., Mohan, E. R., and Ng, B. S. “Development of a hospital mobile platform for logistics tasks”. In: *Digital Communications and Networks* 1.2 (2015), pp. 102–111. ISSN: 23528648. DOI: 10.1016/j.dcan.2015.03.001.

[Aet24] Aethon. *Mobile Robots for Healthcare - Pharmacy, Laboratory, Nutrition and EVS*. 2024. URL: https://aethon.com/mobile-robots-for-healthcare/ (visited on 03/15/2024).

[Ahn+14] Ahn, H. S., Kuo, I.-H., Datta, C., Stafford, R., Kerse, N., Peri, K., Broadbent, E., and MacDonald, B. A. “Design of a Kiosk Type Healthcare Robot System for Older People in Private and Public Places”. In: *Simulation, Modeling, and Programming for Autonomous Robots*. Ed. by Brugali, D., Broenink, J. F., Kroeger, T., and MacDonald, B. A. Vol. 8810. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2014, pp. 578–589. ISBN: 978-3-319-11899-4. DOI: 10.1007/978-3-319-11900-7_49.

[Aik+01] Aiken, L. H., Clarke, S. P., Sloane, D. M., Sochalski, J. A., Busse, R., Clarke, H., Giovannetti, P., Hunt, J., Rafferty, A. M., and Shamian, J. “Nurses’ reports on hospital care in five countries”. In: *Health affairs (Project Hope)* 20.3 (2001), pp. 43–53. DOI: 10.1377/hlthaff.20.3.43.

[Aka24] Akara Robotics. *Akara*. 2024. URL: https://www.akara.ai/ (visited on 03/15/2024).

[AN08] Allison, B. and Nejat, G. “An Expressive Socially Assistive Robot for Health-Care Applications”. In: *Volume 2: 32nd Mechanisms and Robotics Conference, Parts A and B*. ASMEDC, 2008, pp. 755–761. ISBN: 978-0-7918-4326-0. DOI: 10.1115/DETC2008-50082.

[Ama+21] Amato, C., McCanne, L., Yang, C., Ostler, D., Ratib, O., Wilhelm, D., and Bernhard, L. “The hospital of the future: rethinking architectural design to enable new patient-centered treatment concepts”. In: *International journal of computer assisted radiology and surgery* (2021). DOI: 10.1007/s11548-021-02540-9.

[Ama+22] Amato, C., Yang, C., Bernhard, L., Giulianotti, P. C., Kondrat, P., Ratib, O., and Wilhelm, D. "Towards the OR of the future: introducing an adaptive and technology-embracing OR wing layout". In: *International journal of computer assisted radiology and surgery* (2022). DOI: 10.1007/s11548-022-02760-7.

[AV20] Andresen, D. and Voigtländer, T. *Deutscher Herzbericht 2020*. 2020. URL: https://herzstiftung.de/system/files/2021-06/Deutscher-Herzbericht-2020.pdf (visited on 03/15/2024).

[BLX08] Baalbaki, H., Lamiri, M., and Xie, X. "Joint location and configuration of mobile service robots for hospitals". In: *2008 IEEE International Conference on Automation Science and Engineering*. IEEE, 2008, pp. 615–620. ISBN: 978-1-4244-2022-3. DOI: 10.1109/COASE.2008.4626452.

[BX09] Baalbaki, H. and Xie, X. "A decision framework for operation management of reconfigurable mobile service robots in hospitals". In: *IFAC Proceedings Volumes* 42.4 (2009), pp. 151–156. ISSN: 14746670. DOI: 10.3182/20090603-3-RU-2001.0122.

[BX10] Baalbaki, H. and Xie, X. "Distributed rescheduling techniques for a group of mobile reconfigurable service robots". In: (2010).

[BXD10] Baalbaki, H., Xie, X., and Delorme, X. "Mission assignment and scheduling for a team of service robots using evolutionary algorithms". In: *2010 IEEE Workshop on Health Care Management (WHCM)*. IEEE, 2010, pp. 1–6. ISBN: 978-1-4244-4997-2. DOI: 10.1109/WHCM.2010.5441268.

[Bac+17] Bacik, J., Durovsky, F., Biros, M., Kyslan, K., Perdukova, D., and Padmanaban, S. "Pathfinder–Development of Automated Guided Vehicle for Hospital Logistics". In: *IEEE Access* 5 (2017), pp. 26892–26900. DOI: 10.1109/ACCESS.2017.2767899.

[BDO15] Ball, K., Doyle, D., and Oocumma, N. I. "Nursing shortages in the OR: solutions for new models of education". In: *AORN journal* 101.1 (2015), pp. 115–136. DOI: 10.1016/j.aorn.2014.03.015.

[BKW20] Bänziger, T., Kunz, A., and Wegener, K. "Optimizing human–robot task allocation using a simulation tool based on standardized work descriptions". In: *Journal of Intelligent Manufacturing* 31 (2020), pp. 1635–1648.

[Bau+22] Baumann, O., Lenz, A., Hartl, J., Bernhard, L., and Knoll, A. C. "Intuitive teaching of medical device operation to clinical assistance robots". In: *International journal of computer assisted radiology and surgery* (2022). DOI: 10.1007/s11548-022-02802-0.

[BBT01] Bennewitz, M., Burgard, W., and Thrun, S. "Optimizing schedules for prioritized path planning of multi-robot systems". In: *Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation (Cat. No.01CH37164)*. IEEE, 2001, pp. 271–276. ISBN: 0-7803-6576-3. DOI: 10.1109/ROBOT.2001.932565.

[BMC09] Bepko, R. J., Moore, J. R., and Coleman, J. R. "Implementation of a pharmacy automation system (robotics) to ensure medication safety at Norwalk hospital". In: *Quality management in health care* 18.2 (2009), pp. 103–114. ISSN: 1063-8628. DOI: 10.1097/qmh.0b013e3181a02771.

[Ber+07] Berbeglia, G., Cordeau, J.-F., Gribkovskaia, I., and Laporte, G. "Static pickup and delivery problems: a classification scheme and survey". In: *TOP* 15.1 (2007), pp. 1–31. ISSN: 1134-5764. DOI: 10.1007/s11750-007-0009-0.

[Ber+22] Bernhard, L., Amalanesan, A. F., Baumann, O., Rothmeyer, F., Hafner, Y., Berlet, M., Wilhelm, D., and Knoll, A. "Mobile service robots for the operating room wing: balancing cost and performance by optimizing robotic fleet size and composition". In: *International journal of computer assisted radiology and surgery* (2022). DOI: 10.1007/s11548-022-02735-8.

[BKW23] Bernhard, L., Knoll, A., and Wilhelm, D. "Multi-task mission planning for autonomous service robots within the OR wing". In: *CARS 2023—Computer Assisted Radiology and Surgery Proceedings of the 37th International Congress and Exhibition Munich, Germany, June 20–23, 2023* 18.1 (2023), pp. 1–123. DOI: 10.1007/s11548-023-02878-2.

[Ber+20] Bernhard, L., Ostler, D., Feußner, H., and Wilhelm, D. "Integrating autonomously navigating assistance systems into the clinic: guiding principles and the ANTS-OR approach". In: *International journal of computer assisted radiology and surgery* (2020). DOI: 10.1007/s11548-020-02137-8.

[Ber+24] Bernhard, L., Schwingenschlögl, P., Hofmann, J., Wilhelm, D., and Knoll, A. "Boosting the hospital by integrating mobile robotic assistance systems: a comprehensive classification of the risks to be addressed". In: *Autonomous Robots* 48.1 (2024). ISSN: 0929-5593. DOI: 10.1007/s10514-023-10154-0.

[Bil+20] Billings, J., Greene, T., Kember, T., Grey, N., El-Leithy, S., Lee, D., Kennerley, H., Albert, I., Robertson, M., Brewin, C. R., and Bloomfield, M. A. P. "Supporting Hospital Staff During COVID-19: Early Interventions". In: *Occupational medicine (Oxford, England)* 70.5 (2020), pp. 327–329. DOI: 10.1093/occmed/kqaa098.

[Blo11] Bloss, R. "Mobile hospital robots cure numerous logistic needs". In: *Industrial Robot: An International Journal* 38.6 (2011), pp. 567–571. ISSN: 0143-991X. DOI: 10.1108/01439911111179075.

[Bok+20] Bokemeyer, C., Lüftner, D., Hallek, M., and Weißinger, F. *Deutschlandweite Prognose der bevölkerungsbezogenen Morbiditätserwartung für häufige Krebserkrankungen: Auswirkungen auf die Versorgung*. 2020. ISBN: 978-3-9818079-5-0. URL: https://www.dgho.de/publikationen/schriftenreihen/demografischer-wandel/dgho_gpsr_xiv_web.pdf (visited on 03/15/2024).

[Bor17] Borčinova, Z. "Two models of the capacitated vehicle routing problem". In: *Croatian Operational Research Review* 8.2 (2017), pp. 463–469. ISSN: 18480225. DOI: 10.17535/crorr.2017.0029.

[Bro+15] Broadbent, E., Orejana, J. R., Ahn, H. S., Xie, J., Rouse, P., and MacDonald, B. A. "The cost-effectiveness of a robot measuring vital signs in a rural medical practice". In: *2015 24th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 2015, pp. 577–581. ISBN: 978-1-4673-6704-2. DOI: 10.1109/ROMAN.2015.7333668.

[Bue+17] Buechel, M., Hinz, G., Ruehl, F., Schroth, H., Gyoeri, C., and Knoll, A. "Ontology-based traffic scene modeling, traffic regulations dependent situational awareness and decision-making for automated vehicles". In: *2017 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2017, pp. 1471–1476. ISBN: 978-1-5090-4804-5. DOI: 10.1109/IVS.2017.7995917.

[Cac+20] Caccavale, R., Arpenti, P., Paduano, G., Fontanellli, A., Lippiello, V., Villani, L., and Siciliano, B. "A Flexible Robotic Depalletizing System for Supermarket Logistics". In: *IEEE Robotics and Automation Letters* 5.3 (2020), pp. 4471–4476. DOI: 10.1109/LRA.2020.3000427.

[Car+09] Carpintero, E., Perez, C., Morales, R., Garcia, N., Candela, A., and Azorin, J. "Development of a robotic scrub nurse for the operating theatre". In: *2010 3rd IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechatronics*. IEEE, 2009, pp. 504–509. ISBN: 978-1-4244-7708-1. DOI: 10.1109/BIOROB.2010.5626941.

[Car+06] Carreira, F., Canas, T., Silva, A., and Cardeira, C. "i-Merc: A Mobile Robot to Deliver Meals inside Health Services". In: *2006 IEEE Conference on Robotics, Automation and Mechatronics*. IEEE, 2006, pp. 1–8. ISBN: 1-4244-0024-4. DOI: 10.1109/RAMECH.2006.252651.

[Car14] Carus, T. *Operationsatlas Laparoskopische Chirurgie*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014. ISBN: 978-3-642-31245-8. DOI: 10.1007/978-3-642-31246-5.

[CM14] Cepolina, F. E. and Muscolo, G. G. "Design of a robot for hygienization of walls in hospital environments". In: *ISR/Robotik 2014; 41st International Symposium on Robotics*. 2014, pp. 1–7.

[Cha+21] Chai, P. R., Dadabhoy, F. Z., Huang, H.-W., Chu, J. N., Feng, A., Le, H. M., Collins, J., Da Silva, M., Raibert, M., Hur, C., Boyer, E. W., and Traverso, G. "Assessment of the Acceptability and Feasibility of Using Mobile Robotic Systems for Patient Evaluation". In: *JAMA network open* 4.3 (2021), e210667. DOI: 10.1001/jamanetworkopen.2021.0667.

[Cha+23] Chakraa, H., Guérin, F., Leclercq, E., and Lefebvre, D. "Optimization techniques for Multi-Robot Task Allocation problems: Review on the state-of-the-art". In: *Robotics and Autonomous Systems* (2023), p. 104492.

[CK21] Cheikhrouhou, O. and Khoufi, I. "A comprehensive survey on the Multiple Traveling Salesman Problem: Applications, approaches and taxonomy". In: *Computer Science Review* 40 (2021), p. 100369. ISSN: 15740137. DOI: 10.1016/j.cosrev.2021.100369.

[CCK20] Chowdhary, R. R., Chattopadhyay, M. K., and Kamal, R. "Orchestrator Controlled Navigation of Mobile Robots in a Static Environment". In: *Innovations in Electronics and Communication Engineering*. Ed. by Saini, H. S., Singh, R. K., Tariq Beg, M., and Sahambi, J. S. Vol. 107. Lecture Notes in Networks and Systems. Singapore: Springer Singapore, 2020, pp. 193–206. DOI: 10.1007/978-981-15-3172-9_20.

[CV10] Coltin, B. and Veloso, M. “Mobile robot task allocation in hybrid wireless sensor networks”. In: *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2010, pp. 2932–2937. ISBN: 978-1-4244-6674-0. DOI: 10.1109/IROS.2010.5649893.

[CCZ14] Conforti, M., Cornuéjols, G., and Zambelli, G. *Integer programming*. Vol. 271. Graduate texts in mathematics. Cham et al.: Springer, 2014. ISBN: 978-3-319-11008-0. DOI: 10.1007/978-3-319-11008-0.

[Cre+16] Cremer, S., Doelling, K., Lundberg, C. L., McNair, M., Shin, J., and Popa, D. “Application requirements for Robotic Nursing Assistants in hospital environments”. In: *Sensors for Next-Generation Robotics III*. Ed. by Popa, D. and Wijesundara, M. B. J. SPIE Proceedings. SPIE, 2016, 98590E. DOI: 10.1117/12.2229241.

[Cro+18] Croghan, S. M., Carroll, P., Reade, S., Gillis, A. E., and Ridgway, P. F. “Robot Assisted Surgical Ward Rounds: Virtually Always There”. In: *Journal of innovation in health informatics* 25.1 (2018), p. 982. DOI: 10.14236/jhi.v25i1.982.

[Cze+21] Czempiel, T., Paschali, M., Ostler, D., Kim, S. T., Busam, B., and Navab, N. “Opera: Attention-regularized transformers for surgical phase recognition”. In: *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part IV 24*. Springer. 2021, pp. 604–614.

[DS83] Davis, R. and Smith, R. G. “Negotiation as a metaphor for distributed problem solving”. In: *Artificial Intelligence* 20.1 (1983), pp. 63–109. ISSN: 00043702. DOI: 10.1016/0004-3702(83)90015-2.

[DD05] Deary, I. J. and Der, G. “Reaction Time, Age, and Cognitive Ability: Longitudinal Findings from Age 16 to 63 Years in Representative Population Samples”. In: *Aging, Neuropsychology, and Cognition* 12.2 (2005), pp. 187–215. ISSN: 1382-5585. DOI: 10.1080/13825580590969235.

[Dem20] Demaitre, E. *XDBOT developed at NTU Singapore to disinfect large surfaces*. 2020. URL: https://www.therobotreport.com/xdbot-developed-ntu-singapore-disinfect-large-areas-covid-19/ (visited on 03/15/2024).

[Deu22] Deutsche Diabetes Gesellschaft (DDG) and diabetesDE – Deutsche Diabetes-Hilfe. *Deutscher Gesundheitsbericht Diabetes 2022: Die Bestandsaufnahme*. 2022. URL: https://www.deutsche-diabetes-gesellschaft.de/fileadmin/user_upload/Gesundheitsbericht_2022_final.pdf (visited on 03/15/2024).

[Deu20] Deutsche Interdisziplinäre Vereinigung für Intensiv- und Notfallmedizin. *Entscheidungen über die Zuteilung intensivmedizinischer Ressourcen im Kontext der COVID-19-Pandemie - Version 2: Klinisch-ethische Empfehlungen*. 2020.

[Deu21] Deutsches Krankenhaus Institut. *Deutsches Krankenhaus Barometer: Umfrage 2021*. 2021.

[Dil24] Diligent Robotics. *Moxi – Diligent Robotics*. 2024. URL: https://www.diligentrobots.com/moxi (visited on 03/15/2024).

[DIN20] DIN Deutsches Institut für Normung e. V. *Flurförderzeuge – Sicherheitstechnische Anforderungen und Verifizierung (Industrial trucks – Safety requirements and verification): Teil 4: Fahrerlose Flurförderzeuge und ihre Systeme (Part 4: Driverless industrial trucks and their systems)*. Berlin, 2020.

[Din+14] Ding, J., Lim, Y.-J., Solano, M., Shadle, K., Park, C., Lin, C., and Hu, J. "Giving patients a lift - the robotic nursing assistant (RoNA)". In: *2014 IEEE International Conference on Technologies for Practical Robot Applications (TePRA)*. IEEE, 2014, pp. 1–5. ISBN: 978-1-4799-4605-1. DOI: 10.1109/TePRA.2014.6869137.

[Eis+15] Eisenbach, M., Vorndran, A., Sorge, S., and Gross, H.-M. "User recognition for guiding and following people with a mobile robot in a clinical environment". In: *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2015, pp. 3600–3607. ISBN: 978-1-4799-9994-1. DOI: 10.1109/IROS.2015.7353880.

[ek 24] ek robotics. *Automated Transport Systems | ek robotics*. 2024. URL: https://ek-robotics.com/en/solutions/automated-transport-systems/ (visited on 03/15/2024).

[Eri+18] Erickson, Z., Clever, H. M., Turk, G., Liu, C. K., and Kemp, C. C. "Deep Haptic Model Predictive Control for Robot-Assisted Dressing". In: *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 4437–4444. ISBN: 978-1-5386-3081-5. DOI: 10.1109/ICRA.2018.8460656.

[Eva94] Evans, J. M. "HelpMate: an autonomous mobile robot courier for hospitals". In: *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'94)*. IEEE, 1994, pp. 1695–1700. ISBN: 0-7803-1933-8. DOI: 10.1109/IROS.1994.407629.

[Eva+89] Evans, J., Krishnamurthy, B., Pong, W., Croston, R., Weiman, C., and Engelberger, G. "HelpMate$^{\text{TM}}$: A robotic materials transport system". In: *Robotics and Autonomous Systems* 5.3 (1989), pp. 251–256. ISSN: 09218890. DOI: 10.1016/0921-8890(89)90049-3.

[Feu+19] Feussner, H., Wilhelm, D., Navab, N., Knoll, A., and Lüth, T. "Surgineering: a new type of collaboration among surgeons and engineers". In: *International journal of computer assisted radiology and surgery* 14.2 (2019), pp. 187–190. DOI: 10.1007/s11548-018-1893-5.

[Fun+03] Fung, W. K., Leung, Y. Y., Chow, M. K., Liu, Y. H., Xu, Y., Chan, W., Law, T. W., Tso, S. K., and Wang, C. Y. "Development of a hospital service robot for transporting task". In: *IEEE International Conference on Robotics, Intelligent Systems and Signal Processing, 2003. Proceedings. 2003*. IEEE, 2003, pp. 628–633. ISBN: 0-7803-7925-X. DOI: 10.1109/RISSP.2003.1285647.

[Gar+12] Garingo, A., Friedlich, P., Tesoriero, L., Patil, S., Jackson, P., and Seri, I. "The use of mobile robotic telemedicine technology in the neonatal intensive care unit". In: *Journal of perinatology : official journal of the California Perinatal Association* 32.1 (2012), pp. 55–63. DOI: 10.1038/jp.2011.72.

[Gar+16] Garingo, A., Friedlich, P., Chavez, T., Tesoriero, L., Patil, S., Jackson, P., and Seri, I. “"Tele-rounding" with a remotely controlled mobile robot in the neonatal intensive care unit”. In: *Journal of telemedicine and telecare* 22.2 (2016), pp. 132–138. DOI: 10.1177/1357633X15589478.

[Gar+20] Garosi, E., Kalantari, R., Zanjirani Farahani, A., Zuaktafi, M., Hosseinzadeh Roknabadi, E., and Bakhshi, E. “Concerns About Verbal Communication in the Operating Room: A Field Study”. In: *Human factors* 62.6 (2020), pp. 940–953. DOI: 10.1177/0018720819858274.

[GM02] Gerkey, B. P. and Mataric, M. J. “Sold!: auction methods for multirobot coordination”. In: *IEEE Transactions on Robotics and Automation* 18.5 (2002), pp. 758–768. ISSN: 1042-296X. DOI: 10.1109/TRA.2002.803462.

[GM00] Gerkey, B. P. and MatariĆ, M. J. “Murdoch”. In: *Proceedings of the fourth international conference on Autonomous agents - AGENTS ’00*. Ed. by Sierra, C., Gini, M., and Rosenschein, J. S. New York, New York, USA: ACM Press, 2000, pp. 203–204. ISBN: 1581132301. DOI: 10.1145/336595.337369.

[GM04] Gerkey, B. P. and MatariĆ, M. J. “A Formal Analysis and Taxonomy of Task Allocation in Multi-Robot Systems”. In: *The International Journal of Robotics Research* 23.9 (2004), pp. 939–954. ISSN: 0278-3649. DOI: 10.1177/0278364904045564.

[Gro+14] Gross, H.-M., Debes, K., Einhorn, E., Mueller, S., Scheidig, A., Weinrich, C., Bley, A., and Martin, C. “Mobile Robotic Rehabilitation Assistant for walking and orientation training of Stroke Patients: A report on work in progress”. In: *2014 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 2014, pp. 1880–1887. ISBN: 978-1-4799-3840-7. DOI: 10.1109/SMC.2014.6974195.

[Gro+17a] Gross, H.-M., Meyer, S., Scheidig, A., Eisenbach, M., Mueller, S., Trinh, T. Q., Wengefeld, T., Bley, A., Martin, C., and Fricke, C. “Mobile robot companion for walking training of stroke patients in clinical post-stroke rehabilitation”. In: *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 1028–1035. ISBN: 978-1-5090-4633-1. DOI: 10.1109/ICRA.2017.7989124.

[Gro+17b] Gross, H.-M., Scheidig, A., Debes, K., Einhorn, E., Eisenbach, M., Mueller, S., Schmiedel, T., Trinh, T. Q., Weinrich, C., Wengefeld, T., Bley, A., and Martin, C. “ROREAS: robot coach for walking and orientation training in clinical post-stroke rehabilitation—prototype implementation and evaluation in field trials”. In: *Autonomous Robots* 41.3 (2017), pp. 679–698. ISSN: 0929-5593. DOI: 10.1007/s10514-016-9552-6.

[GDK09] Grossmann, F. F., Delport, K., and Keller, D. I. “Emergency Severity Index”. In: *Notfall + Rettungsmedizin* 12.4 (2009), pp. 290–292. ISSN: 1434-6222. DOI: 10.1007/s10049-009-1156-7.

[GGH21] Guettari, M., Gharbi, I., and Hamza, S. “UVC disinfection robot”. In: *Environmental science and pollution research international* 28.30 (2021), pp. 40394–40399. DOI: 10.1007/s11356-020-11184-2.

[Guz+21] Guzmán Ortiz, E., Andres, B., Fraile, F., Poler, R., and Ortiz Bas, Á. "Fleet management system for mobile robots in healthcare environments". In: *Journal of Industrial Engineering and Management* 14.1 (2021), p. 55. ISSN: 2013-8423. DOI: 10.3926/jiem.3284.

[HZC16] Han, S., Zhou, X., and Chen, C. "Path planning for multi-robot systems using PSO and Critical Path Schedule Method". In: *2016 IEEE 13th International Conference on Networking, Sensing, and Control (ICNSC)*. IEEE, 2016, pp. 1–6. ISBN: 978-1-4673-9975-3. DOI: 10.1109/ICNSC.2016.7478999.

[HHS10] Hasan, M. K., Hoque, A. S., and Szecsi, T. "Application of a plug-and-play guidance module for hospital robots". In: *Proceedings of the 2010 International Conference on Industrial Engineering and Operations Management*. 2010, pp. 654–659.

[Hor+11] Horan, B., Najdovski, Z., Black, T., Nahavandi, S., and Crothers, P. "OzTug mobile robot for manufacturing transportation". In: *2011 IEEE International Conference on Systems, Man, and Cybernetics*. IEEE, 2011, pp. 3554–3560. ISBN: 978-1-4577-0653-0. DOI: 10.1109/ICSMC.2011.6084220.

[Hu+11] Hu, J., Edsinger, A., Lim, Y.-J., Donaldson, N., Solano, M., Solochek, A., and Marchessault, R. "An advanced medical robotic system augmenting healthcare capabilities - robotic nursing assistant". In: *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 6264–6269. ISBN: 978-1-61284-386-5. DOI: 10.1109/ICRA.2011.5980213.

[Hua+22] Huang, H.-W., Chen, J., Chai, P. R., Ehmke, C., Rupp, P., Dadabhoy, F. Z., Feng, A., Li, C., Thomas, A. J., Da Silva, M., Boyer, E. W., and Traverso, G. "Mobile Robotic Platform for Contactless Vital Sign Monitoring". In: *Cyborg and bionic systems (Washington, D.C.)* 2022 (2022). DOI: 10.34133/2022/9780497.

[Hus+12] Hussain, A., Rivers, P. A., Glover, S. H., and Fottler, M. D. "Strategies for dealing with future shortages in the nursing workforce: a review". In: *Health services management research* 25.1 (2012), pp. 41–47. DOI: 10.1258/hsmr.2011.011015.

[Ili+14] Ilias, B., Nagarajan, R., Murugappan, M., Helmy, K., Omar, A. S. A., and Rahman, M. A. A. "Hospital nurse following robot: hardware development and sensor integration". In: *International Journal of Medical Engineering and Informatics* 6.1 (2014), p. 1. ISSN: 1755-0653. DOI: 10.1504/IJMEI.2014.058521.

[Int21] International Organization for Standardization. *ISO 8373:2021(en), Robotics — Vocabulary*. 2021. URL: https://www.iso.org/obp/ui/#iso:std:iso:8373:ed-3:v1:en (visited on 03/15/2024).

[IQT21] IQTIG - Institut für Qualitätssicherung und Transparenz im Gesundheitswesen. *Bundesauswertung Cholezystektomie: Erfassungsjahr 2019*. Ed. by IQTIG - Institut für Qualitätssicherung und Transparenz im Gesundheitswesen. 2021. URL: https://iqtig.org/downloads/auswertung/2019/che/DeQS_CHE_2019_BUAW_V01_2021-08-12.pdf (visited on 03/15/2024).

[Jac+12] Jacob, M., Li, Y.-T., Akingba, G., and Wachs, J. P. "Gestonurse: a robotic surgical nurse for handling surgical instruments in the operating room". In: *Journal of robotic surgery* 6.1 (2012), pp. 53–63. ISSN: 1863-2483. DOI: 10.1007/s11701-011-0325-0.

[Jay+16] Jayawardena, C., Kuo, I.-H., Broadbent, E., and MacDonald, B. A. "Socially Assistive Robot HealthBot: Design, Implementation, and Field Trials". In: *IEEE Systems Journal* 10.3 (2016), pp. 1056–1067. ISSN: 1932-8184. DOI: 10.1109/JSYST.2014.2337882.

[JL16a] Jeon, S. and Lee, J. "Multi-robot multi-task allocation for hospital logistics". In: *2016 18th International Conference on Advanced Communication Technology (ICACT)*. IEEE, 2016, pp. 339–341. ISBN: 978-8-9968-6506-3. DOI: 10.1109/ICACT.2016.7423384.

[JL16b] Jeon, S. and Lee, J. "The simulator for performance analysis of multiple robots for hospital delivery". In: *2016 13th International Conference on Ubiquitous Robots and Ambient Intelligence (URAI)*. IEEE, 2016, pp. 746–748. ISBN: 978-1-5090-0821-6. DOI: 10.1109/URAI.2016.7734105.

[JL16c] Jeon, S. and Lee, J. "Vehicle routing problem with pickup and delivery of multiple robots for hospital logistics". In: *2016 16th International Conference on Control, Automation and Systems (ICCAS)*. IEEE, 2016, pp. 1572–1575. ISBN: 978-89-93215-11-3. DOI: 10.1109/ICCAS.2016.7832511.

[JL17a] Jeon, S. and Lee, J. "Framework and Modeling of a Multi-robot Simulator for Hospital Logistics". In: *Robot Intelligence Technology and Applications 4*. Ed. by Kim, J.-H., Karray, F., Jo, J., Sincak, P., and Myung, H. Vol. 447. Advances in Intelligent Systems and Computing. Cham: Springer International Publishing, 2017, pp. 213–219. ISBN: 978-3-319-31291-0. DOI: 10.1007/978-3-319-31293-4_17.

[JL17b] Jeon, S. and Lee, J. "Performance analysis of scheduling multiple robots for hospital logistics". In: *2017 14th International Conference on Ubiquitous Robots and Ambient Intelligence (URAI)*. IEEE, 2017, pp. 937–940. ISBN: 978-1-5090-3056-9. DOI: 10.1109/URAI.2017.7992870.

[JLK17] Jeon, S., Lee, J., and Kim, J. "Multi-robot task allocation for real-time hospital logistics". In: *2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 2017, pp. 2465–2470. ISBN: 978-1-5386-1645-1. DOI: 10.1109/SMC.2017.8122993.

[KP20] Kattepur, A. and Purushotaman, B. "RoboPlanner: a pragmatic task planning framework for autonomous robots". In: *Cognitive Computation and Systems* 2.1 (2020), pp. 12–22. DOI: 10.1049/ccs.2019.0025.

[KHE15] Khamis, A., Hussein, A., and Elmogy, A. "Multi-robot Task Allocation: A Review of the State-of-the-Art". In: *Cooperative Robots and Sensor Networks 2015*. Ed. by Koubâa, A. and Martínez-de Dios, J. Vol. 604. Studies in Computational Intelligence. Cham: Springer International Publishing, 2015, pp. 31–51. ISBN: 978-3-319-18298-8. DOI: 10.1007/978-3-319-18299-5_2.

[Kha13] Khan, M. “The Development of a Mobile Medical Robot Using ER1 Technology”. In: *IEEE Potentials* 32.4 (2013), pp. 34–37. ISSN: 0278-6648. DOI: 10.1109/MPOT.2013.2247096.

[KRL09] Kirschling, T. E., Rough, S. S., and Ludwig, B. C. “Determining the feasibility of robotic courier medication delivery in a hospital setting”. In: *American journal of health-system pharmacy : AJHP : official journal of the American Society of Health-System Pharmacists* 66.19 (2009), pp. 1754–1762. DOI: 10.2146/ajhp080184.

[Kno+04] Knoll, A., Scherer, T., Poggendorf, I., Lütkemeyer, D., and Lehmann, J. “Flexible automation of cell culture and tissue engineering tasks”. In: *Biotechnology progress* 20.6 (2004), pp. 1825–1835. ISSN: 8756-7938. DOI: 10.1021/bp049759v.

[Koc05] Kochan, A. “Scalpel please, robot: Penelope’s debut in the operating room”. In: *Industrial Robot: An International Journal* 32.6 (2005), pp. 449–451.

[Kog+19] Kogkas, A., Ezzat, A., Thakkar, R., Darzi, A., and Mylonas, G. “Free-View, 3D Gaze-Guided Robotic Scrub Nurse”. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*. Ed. by Shen, D., Liu, T., Peters, T. M., Staib, L. H., Essert, C., Zhou, S., Yap, P.-T., and Khan, A. Vol. 11768. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2019, pp. 164–172. ISBN: 978-3-030-32253-3. DOI: 10.1007/978-3-030-32254-0_19.

[Koh+18] Kohn, N., Ostler, D., Koller, S., Marahrens, N., Samm, N., Kranzfelder, M., Vogel, T., Wilhelm, D., and Feußner, H. “Telephone call management in the cognitive operating room”. In: *[13th Russian-German Conference on Biomedical Engineering (RGC), RGC, 2018-05-23 - 2018-05-25, Aachen, Germany]*. 2018. DOI: 10.18154/RWTH-CONV-224970.

[KSD13] Korsah, G. A., Stentz, A., and Dias, M. B. “A comprehensive taxonomy for multi-robot task allocation”. In: *The International Journal of Robotics Research* 32.12 (2013), pp. 1495–1512. ISSN: 0278-3649. DOI: 10.1177/0278364913496484.

[KV07] Korte, B. and Vygen, J. *Combinatorial Optimization: Theory and Algorithms*. 4. 4th ed. 2008. Algorithms and Combinatorics. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007. ISBN: 978-3-540-71844-4.

[KE92] Krishnamurthy, B. and Evans, J. “HelpMate: A robotic courier for hospital use”. In: *[Proceedings] 1992 IEEE International Conference on Systems, Man, and Cybernetics*. IEEE, 1992, pp. 1630–1634. ISBN: 0-7803-0720-8. DOI: 10.1109/ICSMC.1992.271504.

[KSW18] Kumar, B., Sharma, L., and Wu, S.-L. “Job Allocation schemes for Mobile Service Robots in Hospitals”. In: *2018 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. IEEE, 2018, pp. 1323–1326. ISBN: 978-1-5386-5488-0. DOI: 10.1109/BIBM.2018.8621424.

[Lit+63] Little, J. D. C., Murty, K. G., Sweeney, D. W., and Karel, C. “An Algorithm for the Traveling Salesman Problem”. In: *Operations Research* 11.6 (1963), pp. 972–989. ISSN: 0030-364X. DOI: 10.1287/opre.11.6.972.

[Liu+14] Liu, H., Stoll, N., Junginger, S., and Thurow, K. “Mobile robotic transportation in laboratory automation: Multi-robot control, robot-door integration and robot-human interaction”. In: *2014 IEEE International Conference on Robotics and Biomimetics (ROBIO 2014)*. IEEE, 2014, pp. 1033–1038. ISBN: 978-1-4799-7397-2. DOI: 10.1109/ROBIO.2014.7090468.

[Lju+12] Ljungblad, S., Kotrbova, J., Jacobsson, M., Cramer, H., and Niechwiadowicz, K. “Hospital robot at work: Something Alien or an Intelligent Colleague?” In: *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work - CSCW ’12*. Ed. by Poltrock, S., Simone, C., Grudin, J., Mark, G., and Riedl, J. New York, New York, USA: ACM Press, 2012, p. 177. ISBN: 9781450310864. DOI: 10.1145/2145204.2145233.

[Lor+22] Lorca-Oró, C., Vila, J., Pleguezuelos, P., Vergara-Alert, J., Rodon, J., Majó, N., López, S., Segalés, J., Saldaña-Buesa, F., Visa-Boladeras, M., Veà-Baró, A., Campistol, J. M., and Abad, X. “Rapid SARS-CoV-2 Inactivation in a Simulated Hospital Room Using a Mobile and Autonomous Robot Emitting Ultraviolet-C Light”. In: *The Journal of infectious diseases* 225.4 (2022), pp. 587–592. DOI: 10.1093/infdis/jiab551.

[LK16] Luscombe, R. and Kozan, E. “Dynamic resource allocation to improve emergency department efficiency in real time”. In: *European Journal of Operational Research* 255.2 (2016), pp. 593–603. ISSN: 03772217. DOI: 10.1016/j.ejor.2016.05.039.

[Lüt+00] Lütkemeyer, D., Poggendorf, I., Scherer, T., Zhang, J., Knoll, A., and Lehmann, J. “First steps in robot automation of sampling and sample management during cultivation of mammalian cells in pilot scale”. In: *Biotechnology progress* 16.5 (2000), pp. 822–828. ISSN: 8756-7938. DOI: 10.1021/bp0001009.

[LVS16] Lutz, M., Verbeek, C., and Schlegel, C. “Towards a robot fleet for intra-logistic tasks: Combining free robot navigation with multi-robot coordination at bottlenecks”. In: *2016 IEEE 21st International Conference on Emerging Technologies and Factory Automation (ETFA)*. IEEE, 2016, pp. 1–4. ISBN: 978-1-5090-1314-2. DOI: 10.1109/ETFA.2016.7733602.

[MMD21] Ma, A. C., Meng, Z., and Ding, X. “Performance Review of Intelligent Guidance Robot at the Outpatient Clinic Setting”. In: *Cureus* 13.8 (2021), e16840. ISSN: 2168-8184. DOI: 10.7759/cureus.16840.

[Mam+14] Mamun, K. A., Sharma, A., Hoque, A. S. M., and Szecsi, T. “Remote patient physical condition monitoring service module for iWARD hospital robots”. In: *Asia-Pacific World Congress on Computer Science and Engineering*. IEEE, 2014, pp. 1–8. ISBN: 978-1-4799-1955-0. DOI: 10.1109/APWCCSE.2014.7053854.

[Mar+19] Marć, M., Bartosiewicz, A., Burzyńska, J., Chmiel, Z., and Januszewicz, P. “A nursing shortage–a prospect of global and local policies”. In: *International nursing review* 66.1 (2019), pp. 9–16.

[Mat01] Matson, K. “The critical "nurse" in the circulating nurse role”. In: *AORN journal* 73.5 (2001), pp. 971–975. DOI: 10.1016/S0001-2092(06)61750-X.

[Med24a] Medical Expo. *Hospital AMR - QC Bot Open - Vecna Technologies*. 2024. URL: https://www.medicalexpo.com/prod/vecna-technologies/product-105903-697488.html (visited on 03/15/2024).

[Med24b] Medical Expo. *Hospital AMR - QC Lift - Vecna Technologies*. 2024. URL: https://www.medicalexpo.com/prod/vecna-technologies/product-105903-697462.html (visited on 03/15/2024).

[Meg+18] Meghdari, A., Shariati, A., Alemi, M., Vossoughi, G. R., Eydi, A., Ahmadi, E., Mozafari, B., Amoozandeh Nobaveh, A., and Tahami, R. “Arash: A social robot buddy to support children with cancer in a hospital environment”. In: *Proceedings of the Institution of Mechanical Engineers. Part H, Journal of engineering in medicine* 232.6 (2018), pp. 605–618. DOI: 10.1177/0954411918777520.

[MLR19] Mohan, J., Lanka, K., and Rao, A. N. “A Review of Dynamic Job Shop Scheduling Techniques”. In: *Procedia Manufacturing* 30 (2019), pp. 34–39. ISSN: 23519789. DOI: 10.1016/j.promfg.2019.02.006.

[Muk+10] Mukai, T., Hirano, S., Nakashima, H., Kato, Y., Sakaida, Y., Guo, S., and Hosoe, S. “Development of a nursing-care assistant robot RIBA that can lift a human in its arms”. In: *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2010, pp. 5996–6001. ISBN: 978-1-4244-6674-0. DOI: 10.1109/IROS.2010.5651735.

[MBW22] Müller, C., Bernhard, L., and Wilhelm, D. “Usage of a graph database for the selection of sterile items in the OR”. In: *International journal of computer assisted radiology and surgery* (2022). DOI: 10.1007/s11548-022-02795-w.

[Mül+20] Müller, S., Wengefeld, T., Trinh, T. Q., Aganian, D., Eisenbach, M., and Gross, H.-M. “A Multi-Modal Person Perception Framework for Socially Interactive Mobile Service Robots”. In: *Sensors (Basel, Switzerland)* 20.3 (2020). DOI: 10.3390/s20030722.

[Mur+09] Murai, R., Sakai, T., Uematsu, H., Nakajima, H., Mitani, K., and Kitano, H. “Conveyance system using autonomous mobile robots”. In: *2009 IEEE Workshop on Advanced Robotics and its Social Impacts*. IEEE, 2009, pp. 54–59. ISBN: 978-1-4244-4393-2. DOI: 10.1109/ARSO.2009.5587074.

[Nam+22] Nambiappan, H. R., Arboleda, S. A., Lundberg, C. L., Kyrarini, M., Makedon, F., and Gans, N. “MINA: A Robotic Assistant for Hospital Fetching Tasks”. In: *Technologies* 10.2 (2022), p. 41. DOI: 10.3390/technologies10020041.

[NMJ21] Nasirian, B., Mehrandezh, M., and Janabi-Sharifi, F. “Efficient Coverage Path Planning for Mobile Disinfecting Robots Using Graph-Based Representation of Environment”. In: *Frontiers in robotics and AI* 8 (2021), p. 624333. DOI: 10.3389/frobt.2021.624333.

[Nej+07] Nejat, G., Allison, B., Gomez, N., and Rosenfeld, A. “The Design of an Interactive Socially Assistive Robot for Patient Care”. In: *Volume 9: Mechanical Systems and Control, Parts A, B, and C*. ASMEDC, 2007, pp. 821–830. ISBN: 0-7918-4303-3. DOI: 10.1115/IMECE2007-41811.

[NF08] Nejat, G. and Ficocelli, M. "Can I be of assistance? The intelligence behind an assistive robot". In: *2008 IEEE International Conference on Robotics and Automation*. IEEE, 2008, pp. 3564–3569. ISBN: 978-1-4244-1646-2. DOI: 10.1109/ROBOT.2008.4543756.

[NSN09] Nejat, G., Sun, Y., and Nies, M. "Assistive Robots in Health Care Settings". In: *Home Health Care Management & Practice* 21.3 (2009), pp. 177–187. ISSN: 1084-8223. DOI: 10.1177/1084822308325695.

[Neu+22] Neumann, J., Uciteli, A., Meschke, T., Bieck, R., Franke, S., Herre, H., and Neumuth, T. "Ontology-based surgical workflow recognition and prediction". In: *Journal of Biomedical Informatics* 136 (2022), p. 104240.

[New14] News Cision. *RobCab - Robotdalen*. 2014. URL: https://news.cision.com/robotdalen/i/robcab,m4803 (visited on 03/15/2024).

[Ngu+13] Nguyen, H. H., Nguyen, T. N., Clout, R., Gibson, A., and Nguyen, H. T. "Development of an assistive patient mobile system for hospital environments". In: *Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual International Conference* 2013 (2013), pp. 2491–2494. DOI: 10.1109/embc.2013.6610045.

[NSI23] NSI Nursing Solutions. *2023 NSI National Health Care Retention & RN Staffing Report*. 2023. URL: https://www.nsinursingsolutions.com/Documents/Library/NSI_National_Health_Care_Retention_Report.pdf (visited on 03/15/2024).

[Nun+17] Nunes, E., Manner, M., Mitiche, H., and Gini, M. "A taxonomy for task allocation problems with temporal and ordering constraints". In: *Robotics and Autonomous Systems* 90 (2017), pp. 55–70. ISSN: 09218890. DOI: 10.1016/j.robot.2016.10.008.

[Okp22] Okpala, P. "Using technology to ameliorate the effects of healthcare staff shortages during pandemics". In: *International Journal of Healthcare Management* 15.3 (2022), pp. 261–266. ISSN: 2047-9700. DOI: 10.1080/20479700.2020.1871280.

[Ozk+09] Ozkil, A. G., Fan, Z., Dawids, S., Aanes, H., Kristensen, J. K., and Christensen, K. H. "Service robots for hospitals: A case study of transportation tasks in a hospital". In: *2009 IEEE International Conference on Automation and Logistics*. IEEE, 2009, pp. 289–294. ISBN: 978-1-4244-4794-7. DOI: 10.1109/ICAL.2009.5262912.

[PAL22] PAL Robotics. *Fast Deployment of Autonomous Vehicles to Hospital - PAL Robotics*. 2022. URL: https://pal-robotics.com/collaborative-projects/fast-deployment-autonomous-vehicles-hospital/ (visited on 03/15/2024).

[Pan15] Panasonic. *Panasonic Autonomous Delivery Robots - HOSPI - Aid Hospital Operations at Changi General Hospital | Business Solutions | Products & Solutions | Blog Posts | Panasonic Newsroom Global*. 2015. URL: https://news.panasonic.com/global/topics/4923 (visited on 03/15/2024).

[PGP17] Pedan, M., Gregor, M., and Plinta, D. "Implementation of Automated Guided Vehicle System in Healthcare Facility". In: *Procedia Engineering* 192 (2017), pp. 665–670. ISSN: 18777058. DOI: 10.1016/j.proeng.2017.06.115.

[Per+12] Perez-Vidal, C., Carpintero, E., Garcia-Aracil, N., Sabater-Navarro, J. M., Azorin, J. M., Candela, A., and Fernandez, E. "Steps in the development of a robotic scrub nurse". In: *Robotics and Autonomous Systems* 60.6 (2012), pp. 901–911. ISSN: 09218890. DOI: 10.1016/j.robot.2012.01.005.

[Pet+18] Petscharnig, S., Schöffmann, K., Benois-Pineau, J., Chaabouni, S., and Keckstein, J. "Early and Late Fusion of Temporal Information for Classification of Surgical Actions in Laparoscopic Gynecology". In: *2018 IEEE 31st International Symposium on Computer-Based Medical Systems (CBMS)*. 2018, pp. 369–374. DOI: 10.1109/CBMS.2018.00071.

[Pfr+14] Pfrommer, J., Stogl, D., Aleksandrov, K., Schubert, V., and Hein, B. "Modelling and orchestration of service-based manufacturing systems via skills". In: (2014), pp. 1–4. DOI: 10.1109/ETFA.2014.7005285.

[Pri04] Prins, C. "A simple and effective evolutionary algorithm for the vehicle routing problem". In: *Computers & Operations Research* 31.12 (2004), pp. 1985–2002. ISSN: 03050548. DOI: 10.1016/S0305-0548(03)00158-8.

[RGP20] Ramos, T. R. P., Gomes, M. I., and Póvoa, A. P. B. "Multi-depot vehicle routing problem: a comparative study of alternative formulations". In: *International Journal of Logistics Research and Applications* 23.2 (2020), pp. 103–120. ISSN: 1367-5567. DOI: 10.1080/13675567.2019.1630374.

[Rob23] Robotic Automation. *Hospital AGVs (Automated Guided Vehicles) - Robotic Automation*. 2023. URL: https://www.roboticautomation.com.au/ra_solutions/hospital-agvs-automated-guided-vehicles/ (visited on 03/15/2024).

[Röh+21] Röhner, E., Mayfarth, A., Sternitzke, C., Layher, F., Scheidig, A., Groß, H.-M., Matziolis, G., Böhle, S., and Sander, K. "Mobile Robot-Based Gait Training after Total Hip Arthroplasty (THA) Improves Walking in Biomechanical Gait Analysis". In: *Journal of clinical medicine* 10.11 (2021). ISSN: 2077-0383. DOI: 10.3390/jcm10112416.

[RKF98] Rossetti, M. D., Kumar, A., and Felder, R. A. "Mobile robot simulation of clinical laboratory deliveries". In: *1998 Winter Simulation Conference. Proceedings (Cat. No.98CH36274)*. IEEE, 1998, pp. 1415–1422. ISBN: 0-7803-5133-9. DOI: 10.1109/WSC.1998.746010.

[SIM03] Sato, K., Ishii, M., and Madokoro, H. "Testing and evaluation of a patrol robot system for hospitals". In: *Electronics and Communications in Japan (Part III: Fundamental Electronic Science)* 86.12 (2003), pp. 14–26. ISSN: 1042-0967. DOI: 10.1002/ecjc.10139.

[SW17] Savkin, A. V. and Wang, C. "A framework for safe assisted navigation of semi-autonomous vehicles among moving and steady obstacles". In: *Robotica* 35.5 (2017), pp. 981–1005. ISSN: 0263-5747. DOI: 10.1017/S0263574715000922.

[Sch+19] Scheidig, A., Jaeschke, B., Schuetz, B., Trinh, T. Q., Vorndran, A., Mayfarth, A., and Gross, H.-M. "May I Keep an Eye on Your Training? Gait Assessment Assisted by a Mobile Robot". In: *IEEE ... International Conference on Rehabilitation Robotics : [proceedings]* 2019 (2019), pp. 701–708. DOI: 10.1109/ICORR.2019.8779369.

[Sch+03] Scherer, T., Poggendorf, I., Schneider, A., Westhoff, D., Zhang, J., Lutkemeyer, D., Lehmann, J., and Knoll, A. "A service robot for automating the sample management in biotechnological cell cultivations". In: *EFTA 2003. 2003 IEEE Conference on Emerging Technologies and Factory Automation. Proceedings (Cat. No.03TH8696)*. IEEE, 2003, pp. 383–390. ISBN: 0-7803-7937-3. DOI: 10.1109/ETFA.2003.1248725.

[Sch94] Schmitt, L. J. "An empirical study computational study of genetic algorithms to solve order problems: an emphasis on TSP and VRPTC". In: *PhD dissertation, University of Memphis* (1994).

[Sch+20] Schmucker, M., Reiswich, A., Pfeifer, C., Mey, V. de, and Haag, M. "Mobile Robotic Telepresence Between Hospital and School: Lessons Learned". In: *Studies in health technology and informatics* 271 (2020), pp. 256–262. ISSN: 0926-9630. DOI: 10.3233/shti200104.

[See+20] Seenu, N., Kuppan Chetty, R. M., Ramya, M. M., and Janardhanan, M. N. "Review on state-of-the-art dynamic task allocation strategies for multiple-robot systems". In: *Industrial Robot: An International Journal* 47.6 (2020), pp. 929–942. ISSN: 0143-991X. DOI: 10.1108/IR-04-2020-0073.

[She+19] Shen, D., Liu, T., Peters, T. M., Staib, L. H., Essert, C., Zhou, S., Yap, P.-T., and Khan, A., eds. *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2019. ISBN: 978-3-030-32253-3. DOI: 10.1007/978-3-030-32254-0.

[Shi+10] Shi, L., Lin, Wang, M., and Zhang, N. "Mission assignment approach of team service robots based on evolutionary algorithm". In: *2010 International Conference on Educational and Network Technology*. IEEE, 2010, pp. 493–498. ISBN: 978-1-4244-7660-2. DOI: 10.1109/ICENT.2010.5532254.

[SHC04] Shieh, M. Y., Hsieh, J. C., and Cheng, C. P. "Design of an intelligent hospital service robot and its applications". In: *2004 IEEE International Conference on Systems, Man and Cybernetics (IEEE Cat. No.04CH37583)*. IEEE, 2004, pp. 4377–4382. ISBN: 0-7803-8567-5. DOI: 10.1109/ICSMC.2004.1401220.

[SS12] Siewert, J. R. and Stein, H. J. *Chirurgie*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012. ISBN: 978-3-642-11330-7. DOI: 10.1007/978-3-642-11331-4.

[SSS13] Silva, J. B., Santos, C., and Sequeira, J. "Developing a timed navigation architecture for hospital delivery robots". In: *2013 IEEE 3rd Portuguese Meeting in Bioengineering (ENBENG)*. IEEE, 2013, pp. 1–4. ISBN: 978-1-4673-4859-1. DOI: 10.1109/ENBENG.2013.6518412.

[SW19] Sitek, P. and Wikarek, J. "Capacitated vehicle routing problem with pick-up and alternative delivery (CVRPPAD): model and implementation using hybrid approach". In: *Annals of Operations Research* 273.1-2 (2019), pp. 257–277. ISSN: 0254-5330. DOI: 10.1007/s10479-017-2722-x.

[Sit+21] Sitek, P., Wikarek, J., Rutczyńska-Wdowiak, K., Bocewicz, G., and Banaszak, Z. "Optimization of capacitated vehicle routing problem with alternative delivery, pick-up and time windows: A modified hybrid approach". In: *Neurocomputing* 423 (2021), pp. 670–678. ISSN: 09252312. DOI: 10.1016/j.neucom.2020.02.126.

[SKC17] Soares, N., Kay, J. C., and Craven, G. "Mobile Robotic Telepresence Solutions for the Education of Hospitalized Children". In: *Perspectives in health information management* 14.Fall (2017), 1e.

[SKM14] Song, B. D., Kim, J., and Morrison, J. R. "Towards real time scheduling for persistent UAV service: A rolling horizon MILP approach, RHTA and the STAH heuristic". In: *2014 International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 2014, pp. 506–515. ISBN: 978-1-4799-2376-2. DOI: 10.1109/ICUAS.2014.6842292.

[Sta+16] Stauder, R., Ostler, D., Kranzfelder, M., Koller, S., Feußner, H., and Navab, N. *The TUM LapChole dataset for the M2CAI 2016 workflow challenge*. 2016. URL: http://arxiv.org/pdf/1610.09278v2.

[Swi24] Swisslog Healthcare. *Hospital Delivery Robot - Relay - Swisslog Healthcare*. 2024. URL: https://www.swisslog-healthcare.com/en-gb/products/transport/relay (visited on 03/15/2024).

[Tak+12] Takahashi, M., Moriguchi, T., Tanaka, S., Namikawa, H., Shitamoto, H., Nakano, T., Minato, Y., Ihama, T., and Murayama, T. "Development of a Mobile Robot for Transport Application in Hospital". In: *Journal of Robotics and Mechatronics* 24.6 (2012), pp. 1046–1053. ISSN: 0915-3942. DOI: 10.20965/jrm.2012.p1046.

[Tak+09] Takahashi, M., Suzuki, T., Cinquegrani, F., Sorbello, R., and Pagello, E. "A mobile robot for transport applications in hospital domain with safe human detection algorithm". In: *2009 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. IEEE, 2009, pp. 1543–1548. ISBN: 978-1-4244-4774-9. DOI: 10.1109/ROBIO.2009.5420402.

[Tak+10] Takahashi, M., Suzuki, T., Shitamoto, H., Moriguchi, T., and Yoshida, K. "Developing a mobile robot for transport applications in the hospital domain". In: *Robotics and Autonomous Systems* 58.7 (2010), pp. 889–899. ISSN: 09218890. DOI: 10.1016/j.robot.2010.03.010.

[Tak+08] Takashima, K., Nakashima, H., Mukai, T., and Hayashi, S. "Scrub Nurse Robot for Laparoscopic Surgery". In: *Advanced Robotics* 22.13-14 (2008), pp. 1585–1601. ISSN: 0169-1864. DOI: 10.1163/156855308X360884.

[Tas+15] Tasaki, R., Kitazaki, M., Miura, J., and Terashima, K. "Prototype design of medical round supporting robot "Terapio"". In: *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 829–834. ISBN: 978-1-4799-6923-4. DOI: 10.1109/ICRA.2015.7139274.

[Ter+14] Terashima, K., Takenoshita, S., Miura, J., Tasaki, R., Kitazaki, M., Saegusa, R., Miyoshi, T., Uchiyama, N., Sano, S., Satake, J., Ohmura, R., Fukushima, T., Kakihara, K., Kawamura, H., and Takahashi, M. “Medical Round Robot – Terapio –”. In: *Journal of Robotics and Mechatronics* 26.1 (2014), pp. 112–114. ISSN: 0915-3942. DOI: 10.20965/jrm.2014.p0112.

[THB09] Thiel, S., Habe, D., and Block, M. “Co-operative robot teams in a hospital environment”. In: *2009 IEEE International Conference on Intelligent Computing and Intelligent Systems*. IEEE, 2009, pp. 843–847. ISBN: 978-1-4244-4754-1. DOI: 10.1109/ICICISYS.2009.5358271.

[TAV18] Toscano, C., Arrais, R., and Veiga, G. “Enhancement of Industrial Logistic Systems with Semantic 3D Representations for Mobile Manipulators”. In: *ROBOT 2017: Third Iberian Robotics Conference*. Ed. by Ollero, A., Sanfeliu, A., Montano, L., Lau, N., and Cardeira, C. Vol. 694. Advances in Intelligent Systems and Computing. Cham: Springer International Publishing, 2018, pp. 617–628. ISBN: 978-3-319-70835-5. DOI: 10.1007/978-3-319-70836-2_51.

[Twi+17] Twinanda, A. P., Shehata, S., Mutter, D., Marescaux, J., Mathelin, M. d., and Padoy, N. “EndoNet: A Deep Architecture for Recognition Tasks on Laparoscopic Videos”. In: *IEEE transactions on medical imaging* 36.1 (2017), pp. 86–97. DOI: 10.1109/TMI.2016.2593957.

[Twi+19] Twinanda, A. P., Yengera, G., Mutter, D., Marescaux, J., and Padoy, N. “RSDNet: Learning to Predict Remaining Surgery Duration from Laparoscopic Videos Without Manual Annotations”. In: *IEEE Transactions on Medical Imaging* 38.4 (2019), pp. 1069–1078. DOI: 10.1109/TMI.2018.2878055.

[US 24] U.S. Department of Health & Human Services. *START Adult Triage Algorithm - CHEMM*. 2024. URL: https://chemm.hhs.gov/startadult.htm (visited on 03/15/2024).

[UVC23] UVCLight. *UV Disinfection Robots | UVC*. 2023. URL: https://www.uvclight.co.uk/ (visited on 03/15/2024).

[UVD24] UVD Robots. *Advanced UV Disinfection Solutions: Discover UVD Robots*. 2024. URL: https://uvd.blue-ocean-robotics.com/ (visited on 03/15/2024).

[Val+22] Valner, R., Masnavi, H., Rybalskii, I., Põlluäär, R., Kõiv, E., Aabloo, A., Kruusamäe, K., and Singh, A. K. “Scalable and heterogenous mobile robot fleet-based task automation in crowded hospital environments-a field test”. In: *Frontiers in robotics and AI* 9 (2022), p. 922835. DOI: 10.3389/frobt.2022.922835.

[Vas+17] Vasquez, A., Kollmitz, M., Eitel, A., and Burgard, W. “Deep Detection of People and their Mobility Aids for a Hospital Robot”. In: *2017 European Conference on Mobile Robots (ECMR)*. IEEE, 2017, pp. 1–7. ISBN: 978-1-5386-1096-1. DOI: 10.1109/ECMR.2017.8098665.

[Vel+15] Veloso, M., Biswas, J., Coltin, B., and Rosenthal, S. “CoBots: Robust symbiotic autonomous mobile service robots”. In: *Twenty-Fourth International Joint Conference on Artificial Intelligence*. 2015.

[Ver09] Verein Deutscher Ingenieure. *Infrastruktur und periphere Einrichtungen für Fahrerlose Transportsysteme (FTS) - Infrastructure and peripheral installations for Automated Guided Vehicle Systems (AGVS)*. Berlin, 2009.

[Vin02] Vince, A. “A framework for the greedy algorithm”. In: *Discrete Applied Mathematics* 121.1-3 (2002), pp. 247–260. ISSN: 0166218X. DOI: 10.1016/S0166-218X(01)00362-6.

[Vor+18] Vorndran, A., Trinh, T., Mueller, S., Scheidig, A., and Gross, H. “How to Always Keep an Eye on the User with a Mobile Robot?” In: *ISR 2018; 50th International Symposium on Robotics*. 2018, pp. 1–7.

[Wag+22] Wagner, L., Bernhard, L., Fuchtmann, J., Asim Karaoglu, M., Ladikos, A., Feußner, H., and Wilhelm, D. “Integrating 3D cameras into sterile surgical environments: A comparison of different protective materials regarding scan accuracy”. In: *Current Directions in Biomedical Engineering* 8.1 (2022), pp. 25–29. DOI: 10.1515/cdbme-2022-0007.

[Wag+23a] Wagner, L., Kolb, S., Leuchtenberger, P., Bernhard, L., Jell, A., and Wilhelm, D. “Deep learning assisted intraoperative instrument cleaning station for robotic scrub nurse systems”. In: *at-Automatisierungstechnik* 71.7 (2023), pp. 572–579.

[Wag+23b] Wagner, L., Kolb, S., Looschen, C., Bernhard, L., Fuchtmann, J., Berlet, M., Fottner, J., Knoll, A., and Wilhelm, D. “Versatile end effector for laparoscopic robotic scrub nurse”. In: *International journal of computer assisted radiology and surgery* (2023). DOI: 10.1007/s11548-023-02892-4.

[Wan+16] Wang, C., Matveev, A. S., Savkin, A. V., Clout, R., and Nguyen, H. T. “A semi-autonomous motorized mobile hospital bed for safe transportation of head injury patients in dynamic hospital environments without bed switching”. In: *Robotica* 34.8 (2016), pp. 1880–1897. ISSN: 0263-5747. DOI: 10.1017/S0263574714002641.

[Wan+15] Wang, C., Savkin, A. V., Clout, R., and Nguyen, H. T. “An Intelligent Robotic Hospital Bed for Safe Transportation of Critical Neurosurgery Patients Along Crowded Hospital Corridors”. In: *IEEE transactions on neural systems and rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society* 23.5 (2015), pp. 744–754. DOI: 10.1109/tnsre.2014.2347377.

[Wan+09] Wang, J., Wei, B.-y., Zhang, Y., and Chen, H. “Design of an autonomous mobile robot for hospital”. In: *2009 IEEE International Symposium on IT in Medicine & Education*. IEEE, 2009, pp. 1181–1186. ISBN: 978-1-4244-3928-7. DOI: 10.1109/ITIME.2009.5236275.

[Wan+21] Wang, Z., Deligianni, F., Voiculescu, I., and Yang, G.-Z. “A Single RGB Camera Based Gait Analysis With A Mobile Tele-Robot For Healthcare”. In: *Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual International Conference* 2021 (2021), pp. 6933–6936. DOI: 10.1109/embc46164.2021.9630765.

[WN14] Wassan, N. and Nagy, G. "Vehicle Routing Problem with Deliveries and Pickups: Modelling Issues and Meta-heuristics Solution Approaches". In: *International Journal of Transportation* 2.1 (2014), pp. 95–110. ISSN: 22877940. DOI: 10.14257/ijt.2014.2.1.06.

[WC00] Weinberg, A. and Creed, F. "Stress and psychiatric disorder in healthcare professionals and hospital staff". In: *Lancet (London, England)* 355.9203 (2000), pp. 533–537. ISSN: 0140-6736. DOI: 10.1016/S0140-6736(99)07366-3.

[Wen+16] Wengefeld, T., Eisenbach, M., Trinh, T. Q., and Gross, H. "May I be your Personal Coach? Bringing Together Person Tracking and Visual Re-identification on a Mobile Robot". In: *Proceedings of ISR 2016: 47st International Symposium on Robotics*. 2016, pp. 1–8.

[Wes+04] Westhoff, D., Stanek, H., Scherer, T., Zhang, J., and Knoll, A. "A flexible framework for task-oriented programming of service robots". In: *Robotik 2004, VDI/VDE-Gesellschaft Mess-und Automatisierungstechnik, VDI-Berichte* (2004).

[Woj+07] Wojtczyk, M., Heidemann, R., Joeris, K., Zhang, C., Burnett, M., Knoll, A., and Konstantinov, K. "The Use of a Mobile Robot for Complete Sample Management in a Cell Culture Pilot Plant". In: *Cell Technology for Cell Products*. Ed. by Smith, R. Dordrecht: Springer Netherlands, 2007, pp. 543–547. ISBN: 978-1-4020-5475-4. DOI: 10.1007/978-1-4020-5476-1_93.

[Wor16] World Health Organization. *Global strategy on human resources for health: Workforce 2030*. 2016.

[WZY16] Wu, Q., Zhao, L., and Ye, X.-C. "Shortage of healthcare professionals in China". In: *BMJ (Clinical research ed.)* 354 (2016), p. i4860. DOI: 10.1136/bmj.i4860.

[Xen23] Xenex Microbial Reduction. *The First and Only FDA Authorized Microbial Reduction Robot for the Healthcare Environment*. 2023. URL: https://xenex.com/ (visited on 03/15/2024).

[Xia+13] Xiao, J., Aydt, H., Lees, M., and Knoll, A. "A partition based match making algorithm for taxi sharing". In: *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*. IEEE, 2013, pp. 2393–2400. ISBN: 978-1-4799-2914-6. DOI: 10.1109/ITSC.2013.6728585.

[Yos+10] Yoshimitsu, K., Masamune, K., Iseki, H., Fukui, Y., Hashimoto, D., and Miyawaki, F. "Development of Scrub Nurse Robot (SNR) systems for endoscopic and laparoscopic surgery". In: *2010 International Symposium on Micro-NanoMechatronics and Human Science*. IEEE, 2010, pp. 83–88. ISBN: 978-1-4244-7995-5. DOI: 10.1109/MHS.2010.5669574.

[Yuk+04] Yukawa, T., Saito, N., Matsuoka, W., Kanda, A., and Okano, H. "Autonomous control of care and welfare robot". In: *IEEE Conference on Robotics and Automation, 2004. TExCRA Technical Exhibition Based*. IEEE, 2004, pp. 53–54. ISBN: 0-7803-8564-0. DOI: 10.1109/TEXCRA.2004.1424993.

[ZAF16] Zander, S., Ahmed, N., and Frank, M. T. "A Survey about the Usage of Semantic Technologies for the Description of Robotic Components and Capabilities". In: *SAMI@ iKNOW*. 2016.

[ZT10] Zilles, K. and Tillmann, B. N. *Anatomie*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010. ISBN: 978-3-540-69481-6. DOI: 10.1007/978-3-540-69483-0.

[Zuk+18] Zukowski, M., Matus, K., Pawluczuk, E., Kondratiuk, M., and Ambroziak, L. "Patients temperature measurement system for medical robotic assistant". In: AIP Conference Proceedings. Author(s), 2018, p. 020084. DOI: 10.1063/1.5066546.