Situating Robots in the Emergency Department

Angelique Taylor, Sachiko Matsumoto, and Laurel D. Riek
Computer Science and Engineering UC San Diego

Abstract

The emergency department (ED) is a safety-critical environment in which mistakes can be deadly and providers are overburdened. Well-designed and contextualized robots could be an asset in the ED by relieving providers of non-value added tasks and enabling them to spend more time on patient care. To support future work in this application domain, in this paper, we characterize ED staff workflow and patient experience, and identify key considerations for robots in the ED, including safety, physical and behavioral attributes, usability, and training. Then, we discuss the task representation and data needed to situate the robot in the ED, based on this domain knowledge. To the best of our knowledge, this is the first work on robot design for the ED that explicitly takes task acuity into account. This is an exciting area of research and we hope our work inspires further exploration into this problem domain.

Introduction

The emergency department (ED) is a fast-paced, safety-critical environment where patients frequently have high levels of acuity [28, 20, 1]. ED clinicians are responsible for many concurrent tasks such as administrative work, diagnosis, and management of complex cases, as well as teaching and liaising with law enforcement, ambulances, and patients' relatives [1, 58]. Because they have so many responsibilities, clinicians are constantly interrupted and have to switch between tasks as a result of unorganized, unplanned, and unpredictable environmental conditions. These conditions often lead to mistakes and clinician burnout which can negatively impact patient outcomes, causing patient safety risks and potentially death [9].

Given the number of critical patient-care decisions that are made in EDs, there is great interest in developing methods to assist providers and improve patient outcomes [55]. Some approaches focus on procedural changes, like having nurses wear vests when performing certain tasks that indicate other people should not interrupt them [63], or physical changes to the built environment. Other approaches , such as the ontology proposed by Tao et al. [54], introduce new ways to represent tasks in the ED.

Figure 1: This photo shows the safety-critical, chaotic conditions that can occur in the emergency department.

Robots are one technology that can potentially improve conditions in the ED. Commercially available robots already work in hospitals to help clinicians deliver and stock materials, clean and sanitize rooms after procedures, and lift patients [18, 35, 64, 14, 40]. Robots also serve as receptionists, assist with rehabilitation for people with dementia and autism, support teams in surgery, and support remote care via telepresent robots [2, 60, 13].

Despite these recent efforts in deploying robots in healthcare settings [48], it will be difficult for robots to execute the simplest tasks in the ED, which is a more challenging environment than in-patient units. The ED is more crowded, chaotic, and has patients with higher levels of acuity. Furthermore, introducing robots into the ED could be disruptive and exacerbate its chaotic nature. As a result, researchers need to carefully consider the existing dynamics of the ED to avoid creating additional layers of complexity.

Despite these challenges, the ED is an exciting area in which robots could make a significant impact. They can free up time of skilled workers so they can focus on the tasks that they are skilled at (complex problem solving, dextrous manipulation, customer service, etc), as opposed to what are referred to as "non-value added tasks," such as material delivery and retrieving samples. Patients' experiences may also be positively impacted by robots, for example, as serving as bedside assistants, educators, or companions [26, 34, 61, 15].

In our work, we are interested in learning how to place robots in the ED, identifying the tasks they could perform, and describing the ED-specific considerations for robotic-assisted ED tasks. We provide insights into the current state of the ED in terms of care delivery from the perspectives of both patients and providers. Using this domain knowledge, we highlight key concerns and contextual considerations for situating robots in the ED. Then, we provide a case study to illustrate how we took these factors into account in our recent work designing a delivery robot for the ED.

To the best of our knowledge, this is the first work to design robots to work in the ED while taking task priority into account. Ultimately, we hope the results of our work will prove useful to robotics researchers working in healthcare as well as other safety critical domains, including search and rescue, first response, and defense.

Characterizing Emergency Department Care Delivery

In this section, we characterize care delivery within the ED. We note that this work is predominantly based on research from hospital emergency departments in the US, and other countries and cultures may have different practices.

Patient Experience

Figure 2 shows a typical trajectory for a patient in most EDs [54]. When patients arrive, they check-in and triage staff (often a nurse) conducts an assessment. Triage staff will take the patient's vital signs, such as heart rate and blood pressure, take their medical history, and record the reason they are visiting the ED. Based on this information, triage staff will determine the Emergency Severity Index (ESI) category of the patient. This is a number from 1 to 5, with 1 being the most acute and 5 being the least acute. For example, an ESI of 1 indicates that a patient requires immediate, life-saving intervention (e.g., a gunshot wound or heart attack), whereas an ESI of 5 might be a head cold.

After admission to the ED, patients receive a more thorough assessment by a nurse, followed by a physician. From this examination, the physician determines what diagnostic tests or treatments are needed. For instance, a patient may need an X-ray or blood test to assist in diagnosis. Treatment may involve bedside procedures, medication, and possibly admittance from the ED to another unit in the hospital that can provide a more appropriate level of care for the patient's condition.

If the patient does not need to be transferred, they will be discharged after treatment. The physician will prescribe any necessary outpatient medications and sign to discharge the patient. A nurse will provide the patient with instructions concerning any medications or further care required. Then the patient will sign any necessary paperwork and depart the ED.

Emergency Department Staff Workflow

Now that we have discussed the ED staff roles and their tasks, we provide a discussion on ED workflow, particularly after the patient has been admitted. Before the provider conducts their assessment, the physician reviews the patient's record as well as the information recorded in triage. Then, the physician goes to the patient's room and conducts a provider assessment where they interview the patient in an effort to collect data for a diagnosis.

Next, the physician performs a physical exam followed by a more in-depth review of the patient's history as well as the relevant medical literature for the patient's hypothesized condition. Then, the physician orders diagnostic tests based on this information. After gathering as much patient data as possible in the time-frame appropriate for the patient's level of acuity, the physician makes a diagnosis and decides on a treatment plan.

The ED is composed of clinical staff, technicians, and administrative staff [54]. The clinical staff are responsible for diagnosing and treating patients, and consists of nurse practitioners, physician assistants, attending physicians, and resident physicians. ED technicians assist clinical staff to provide patient care such as assessments, transportation, and basic procedures. The administrative staff collect and record data about patients, manage billing, and oversee administrative processes in the ED.

Many researchers have conducted ethnographic studies to understand the ED workflow [1, 16, 24, 37]. These studies identified many tasks that ED staff perform which include direct and indirect patient care. Direct patient care tasks are conducted at the patient's bedside, and include checking a patient's medical history, performing procedures, and communicating with the patient and their family. ED staff can spend from 25-40% of their time on direct patient care [1].

Indirect patient care tasks include charting (documentation), ordering diagnostic tests and medication, communicating with other ED staff, procedural planning, and teaching. Other activities performed in the ED include administrative tasks such as meetings, writing reports and emails, and staffing. Other tasks include research activities, and educational activities such as professional development and reading. ED staff can spend from 45-65% of their time on indirect patient care [1].

Situating Robots in the ED

The ED is a non-deterministic environment where staff have a high workload, are under time-sensitive constraints, and must frequently make decisions under uncertainty [20]. No robots, to the best of our knowledge, work in this space today. As such, we provide a discussion on designing robots for the ED using the domain knowledge from Sections - to inform our discussion.

Key Considerations for Robots in the ED

When considering developing robots for the ED, it is important to consider five factors: the robot's safety, its physical and behavioral attributes, its acceptability, and the training required to situate it within the ED.

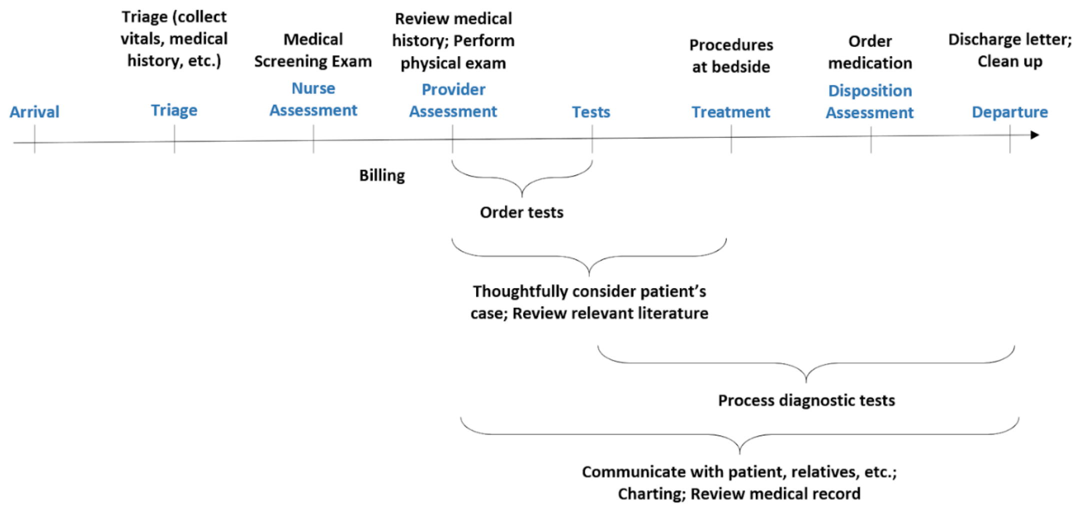

Figure 2: A timeline of a typical patient visit to the ED [54]. The blue text shows the patient's trajectory, and the black text shows the tasks ED staff perform.

Safety. If robots are introduced to the ED, they must be built to be safe around people. For instance, the robot must be able to avoid colliding with people. This could be difficult given that ED hallways are often crowded, and people often move quickly from place to place due to the demands of their work. Additionally, the robot's joints should be compliant, so that even if someone collides with it or the robot makes an error, the person is less likely to get hurt.

In addition to these constraints, which apply generally when people work in close proximity with robots [38, 57], there are also safety concerns specific to the ED. For instance, the robot should be easy to clean and disinfect, so that it does not spread infection. Furthermore, patients are often attached to machines via wires and tubes, such as IV drips, blood oxygen saturation monitors, and dialysis machines. It is critical that robots do not run over these tubes or dislodge them from patients. If the robot dislodged an oxygen monitor, this might merely be an inconvenience, but interfering in a process like dialysis could be deadly. Furthermore, many procedures in the ED are critical to a patient's survival, and robots must be aware enough of their environment to not interrupt these procedures.

Physical Attributes. A robot designed for the ED may have a variety of physical attributes. The robot could be mobile or stationary, and could have no manipulators or several. These attributes will inform the capabilities of the robot. For instance, a stationary robot would not be able to deliver supplies, while it might be unnecessary for a patient bedside robot to be mobile. Additionally, the robot could have a humanoid, zoomorphic, or mechanistic morphology, which will affect people's perception of it and their expectations around it [47, 48]. This can contribute to the acceptability of the robot within the ED, which in turn will affect its functionality and usefulness.

Behavioral Attributes. In addition to physical attributes, behavioral attributes will also affect a robot's performance and acceptance. For instance, should the robot's behavior be goal and performance-driven, or should it try to behave socially? This likely will depend on the role of the robot. For example, a bedside robot would likely need to engage in social interaction, while it might be better for a delivery robot to be performance-driven, such as ensuring that it takes the quickest path while delivering supplies. However, it may still be beneficial for a performance-driven robot to take some social context into account. A mobile robot, for instance, might cause less disturbance and better avoid people if it conforms to social norms as it moves [39, 50, 59, 41].

Furthermore, it is important to consider how the robot will exchange information with people. For instance, providers must be able to instruct the robot to do tasks, and the robot must be able to inform people it is unavailable to carry out their request because if it is already occupied with another task. The robot should support multiple modes of communication both to support people with disabilities as well as to be adaptive to a range of different hospital contexts, which can vary widely [21]. For example, it could have a touch screen that people use to assign it tasks and that displays its state visually, or it could communicate verbally. The robot could also use gestures and implicit communication modalities to more clearly inform people of its state [52] and better understand people giving it instructions.

Acceptability. Because mistakes in EDs can cause grave harm or death, it is important that the introduction of the robot does not cause excessive disturbance or hinder clinicians. To ensure that the robot fills a need for clinicians and does not disrupt the ED, researchers should closely collaborate with ED staff. Researchers can interview providers and perform careful contextual analysis on the transcripts to identify themes across different people. They can also directly involve providers in the design of the robot.

Subsequently, the robot must be thoroughly tested and possible outcomes fully considered before it is placed in the ED. This should include not only testing the technical aspects of the robot, but also consulting with providers in the

ED to determine how the robot can best fill providers' needs.

Furthermore, research teams should conduct longitudinal assessments after a robot has been deployed in the ED to ensure it remains useful and does not impede clinicians in their work. Researchers should analyze how the robot affects workflow in the ED, care delivery, and patient experience. This should also include an assessment of ED staff members' perceptions of the robot.

Longitudinal studies are also important because people will habituate to the robot. People might use the robot initially because it is novel but then stop using it if does not make their tasks easier. On the other hand, people may initially be wary of using the robot but could come to use it frequently if it significantly eases their workload. For the robot to be useful long term, it will need to be designed well, so it is intuitive and functional to facilitate clinicians' work.

Training. In some situations, it may be essential to require all staff members to undergo extensive training to use the robot. This will reduce mistakes and enable providers to perform more complex tasks with the robot. Additionally, people working around the robot may also need training to be safe, even if they do not directly interact with the robot. Therefore, the type of training each person receives should be tailored to their role with the robot.

Additionally, for certain tasks, it might be preferable to require little to no training to work with the robot. For instance, a provider should not need to train a patient to use a bedside robot, as this would add to the workload of ED staff members. However, the robot itself could train patients by walking them through a tutorial. Thus, the type of user and task context affect the amount of required training to use the robot.

Framing Robots for the Emergency Department

Task Representation. For the past three years, we have been engaged in close collaboration with ED staff at our institution's academic medical center to consider how robots may be beneficial to their work practices and to the ED patient experience. Our colleagues identified several problems, including: long wait times, slow triage check-in, collecting and sending samples to the lab, interruptions in patient care, and delivering supplies in time-sensitive situations. We found these problems intriguing, and have been exploring how intelligent robotic systems might be potentially useful.

In order to develop robots for a new problem domain such as the ED, we need to understand what tasks ED staff perform and how to represent them. We identified several tasks that ED staff perform in Section . In order to develop a task representation which might be suitable for robots to solve, we need a clear understanding of the problems or bottlenecks that occur when ED staff perform these tasks. For example, when clinicians work at a patient's bedside, they are often interrupted, which leads to inaccurate/incomplete documentation, and degraded patient care. A robot working in the ED should understand this and be wary of interrupting a clinician working at a patient's bedside. However, if there is an emergency, it might need to interrupt the clinician. Therefore, it needs a well-defined representation of the relative priorities of the tasks involved. With concrete understanding of an ED task, we can represent its complexity and begin to examine its nuances, and explore how a robot may be able (or may not be able to) support care delivery.

Another aspect of task representation to consider are the responsibilities or goals of the robot when it works alongside clinical teams. In recent work, we worked with nurses to co-design robots to support and empower them [55]. Through our collaborations with several medical centers, we found that nurses are often penalized for speaking up when identifying mistakes made by others in clinical teams. As the primary patient advocate, nurses are uniquely positioned to stop behaviors that lead to safety risks. Thus, we discovered the need to support and empower nurses.

A similar approach can be used to design robots for the ED. Researchers can collaborate with hospital EDs to design robots that better address the needs of providers. However, as our recent work suggests, collaboration with people in different roles in the ED can result in different priorities being reflected in the robot design. For example, the goals of nurses are different from the goals of physicians, so nurses will likely design robots to fulfill their unique goals. Therefore, identifying the user group for robots is an important part of this work.

Task temporality is critical in the ED due to the time-sensitive nature of acute patients. One way to represent task temporally is using time-motion analysis [29]. This is a commonly used technique in ethnographic research conducted in healthcare [65, 22], which can provide a roadmap for technology researchers interested in working within a healthcare context . Though time-motion analysis, clinical breakdowns can be broken down by the hour, day, or even months in order to identify its long-term causes and effects.

Another important aspect of task representation is patient acuity. For high acuity patients, it is essential that the robot act quickly and accurately. High acuity patients have time-sensitive conditions and therefore have a high priority when they enter the ED. At any moment in time, the ED can admit a high acuity patient that requires immediate attention from a physician working with low acuity patients. Yet EDs typically have intra-day fluctuations in patient census, and it can be challenging to prioritize treatment of newly admitted, high-acuity patients during high census periods. This must be taken into account by ED staff and robots alike.

With an adequate understanding of the problem space, another important aspect of designing robots for the ED is to understand what data a robot requires to accomplish its goals.

Information Gathering. In the ED, there is the potential to collect data from sensors mounted on the robot, sensors worn by a patient on ED staff members, external sensors in the environment, and the patient's electronic medical record. In our work we focus on the first two topics, though see [7, 25, 3] for surveys on the others.

Robots typically use visual sensors such as RGB, RGB-D (color and depth), and LiDAR. These sensors provide data that can be used in vision and learning algorithms that enable robots to understand the current state of the environment. For example, a large body of work addresses activity recognition using RGB data [49, 66], which can use useful for

Figure 3: These figures show the hallways of emergency departments. Patients are often treated in hallways when the ED bedrooms are full. Placing patients in hallways is a way to handle an overflow of patients. The hallways are often cluttered, over crowded, and clinicians treating time-sensitive safety-critical patients.

developing robots to perceive the actions of clinical workers. Activity recognition has been approached using machine learning such as supervised learning algorithms based on human joint positions [43, 51]. More recently, researchers have employ deep learning techniques such as recurrent neural networks to address this problem [5].

For sensors that go on robots, another consideration is where to place the sensor(s). For the purposes of this paper, we consider sensors placed onboard the robot, which is known as an ego-centric (or first-person) perspective. Typically, LiDAR is placed at the robot's base to be used for Simultaneous Localization and Mapping (SLAM), a commonly used technique for navigation [6]. Visual sensors, like RGB-D, are typically placed at human height on robots that interact with people [56, 42]. This positioning provides an adequate field-of-view for robots to observation its teammates.

An alternative to visual sensors are non-visual sensors that can be placed on the human body, in our case, ED staff or patients. Common non-visual, non-intrusive sensors include internal measurement units (IMUs), which measure accelerations and velocities; sensors that measure physiological information about the body, such as electromyography (EMG), which detect the electrical activity of muscles; and Radio Frequency Identification (RFID) tags that provide positional information (such as where a staff member, patient, or piece of equipment is located) These systems avoid many of the data privacy concerns that occur with video data, which is a common concern in healthcare settings [32, 33]. Non-intrusive sensors are also used to avoid occlusions and poor image quality, issues relevant to visual sensors.

However, these sensors often communicate with each other and the robot via WiFi or bluetooth. Reliable, fast WiFi may be limited in the ED, and thick walls can block both WiFi and bluetooth signals. Additionally, non-visual, wearable sensors are limited to the person or piece of equipment they are on, whereas visual sensors can observe many different entities at once.

Case Study: Delivery Robots

Based on our aforementioned work with ED staff, we have been collaboratively designing a robot to deliver materials. One major topic of concern was the need to fetch equipment quickly. When clinicians need equipment for a patient with high acuity, a member of the team needs to find and fetch equipment, leaving the clinical team short one person. If a robot could deliver equipment and materials, the aforementioned delivery person could instead remain focused on direct patient care tasks .

However, this is a challenging problem for robots. For one thing, hallways in EDs can be very chaotic. Because there are a limited number of rooms, patients are often treated in the hallways. If a robot interrupts providers performing a life-saving treatment on a patient, it could result in the patient dying. On the other hand, it may be acceptable for the robot to interrupt teams performing lower priority tasks. For instance, the robot might move through a group of conversing clinicians if doing so allows it to more quickly complete its delivery.

Thus, when delivering supplies, the robot must account for the priority of the task clinicians are performing. In our work, we formulate this research as a socially-aware path planning problem, which we describe below in a case study. We present a simple algorithm and task representation, derived from our co-design activities with ED staff and on our insights derived from the literature.

Task Representation

Reinforcement learning (RL) is a popular technique in robotics and is commonly used for path planning [30, 31, 10, 11, 17, 45, 53, 27, 8, 36, 44]. It provides a framework where an agent explores or exploits an environment through exploration. Recently, researchers have combined deep learning methods with RL [27, 8]. However, prior RL approaches do not address our problem as we need to account to various levels of patient acuity, whether a hallway is too cluttered or chaotic for a robot to navigate through, and how to handle situations when all hallways in the ED are cluttered and the robot still needs to deliver materials to a patient with a life-threatening condition.

Nevertheless, RL is particularly well-suited for this problem because the environment can be represented as a Markov Decision Process (MDP), where we can easily represent our scenario in terms of the environment, actions of the robot, and its goals. MDPs are particularly useful for learning decision making policies in uncertain environments. The robot learns to plan paths through exploration of the environment in an unsupervised manner through a penalty-reward system.

Figure 4 shows a simplified example of the scenario we envision. In this scenario, a provider (green) needs supplies and asks the robot to deliver them. The robot (blue) must plan a path from its current location to the provider. Two potential paths are approximately equal in length. However, along one of these paths, clinicians are treating a high-priority patient (orange). Along the other, clinicians are engaged in a lower priority task, such as conversing (brown).

There are no clinicians along a third path, but it is much longer than the other two.

In this scenario, the robot should not choose the path with the high-priority team. However, it might be acceptable for it to choose the path with the low-priority team, even though it will interrupt clinicians.

The goal of our robot is to plan a path from a starting position to a goal position while avoiding hallways where high acuity patients are being treated (to the best of its ability). The robot's behavior is generated by a policy π which maps states to actions that maximizes its overall reward R_t where γ is the discounted factor and r_t is the reward at time t (See Equ. 1).

$$R_t = \sum_{t=0}^{\infty} \gamma^t r_t \tag{1}$$

The agent's behavior is formalized by a policy π that maps states S to a set of actions A where,

- States $S = \{s_1, s_2, \dots, N\}$ are locations on the map.
- Actions $A = \{a_1, a_2, \dots, M\}$ are a move from one node of the graph node to another node.
- Rewards R encodes the level of priority of a clinical team.

We use an action-value function $Q^*(s, a)$ to determine the value of a given state. By maximizing that action-value function, we maximize the expected rewards over all a series of actions following a policy π.

$$Q^*(s_t, a_t) = \max_{\pi} \mathbf{E}[R_t | s_t, a_t, \pi] \tag{2}$$

$$L_i(\theta_i) \leftarrow \mathbf{E}[(r + \gamma \max_{a}^{\prime} Q(s', a'; \theta_i^-) - Q(s, a; \theta_i))^2] \tag{3}$$

To maximize the rewards over all actions, the optimal action-value function obeys the Bellman Equation shown below. As done commonly in the literature [44, 23, 12, 4], we use Q-Learning [62], a model-free algorithm to teach that robot a policy of what actions to take under certain penalties and rewards.

$$Q^*(s_t, a_t) = \mathbf{E}_{s_t \sim S}[r_t + \gamma \max_{a_t} Q^*(s_t, a_t) | s, a] \tag{4}$$

The inputs to our RL algorithm are a map of the environment and the location of the robot, clinical teams and their priority, and the user (See Fig. 4). To represent the environment of the ED, we use a topological graph overlaid on our map of the environment. The graph nodes represent waypoints throughout the ED such as the location of the robot (blue), groups of clinicians talking in the hallway (yellow), groups of clinicians working on life-critical patients (orange), and the user that made a request for delivery (green), as well as intersections in the hallways (black). The robot's position (blue) is the starting position and the user's position (green) is the goal location.

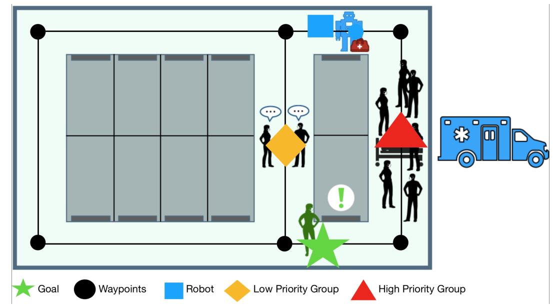

Figure 4: This is a simple example of an ED environment. We model the environment with a graph. The robot (blue) needs to deliver supplies to waypoints in the graph (black nodes). It must navigate around both low priority groups (brown) and high priority groups (orange). The ED staff member who made the original request appears in green. The goal of the delivery robot is to generate a path from its current position to the ED staff requester without interrupting the safety-critical team (Figure inspired by [19]).

Scenarios in the Emergency Department

We present three scenarios in the ED for our agent to learn how to plan paths (See Figure 5). The level of difficulty of the scenarios range from easy, medium, and hard in terms of the number of clinical teams that the agent needs to consider for its paths.

We trained a Q-Learning algorithm [62] over 700 episodes. Our agent takes an action using a ϵ-greedy strategy as commonly done in the literature [46]. This strategy explores a policy by choosing a random action with probability $\in$ [0,1] We incorporate the patient priority in our reward function as shown below. We use the following parameters to train a Q-Learning algorithm: $r_h = -1$, $r_h = -5$, $r_h = 100$ and d is the length of the path represented between two states being considering. A negative reward means to avoid a hallway in the ED and a positive reward represents the location of the goal. We use a discount coefficient $\gamma = 0.8$

$$r = \begin{cases} d + r_h, & \text{if high priority group} \\ d + r_l, & \text{if low priority group} \\ r_g, & \text{if found goal} \end{cases}$$

The resulting paths for our scenarios are shown in Figure 5. Our agent generated easy and medium as expected. In the hard scenario, the agent generates behavior that depicts the key challenge in our current work – that is, how can an agent generate paths when the ED is when all paths have patients with various levels of acuity? When all paths contain groups, both low priority and high priority, the agent never finds a solution because the desired behavior must be captured by a more complex policy. This forms the basis of our existing work as we design delivery robots for the ED under conditions where all hallways can be cluttered and occupied by clinical teams that perform procedures on highly acute patients.

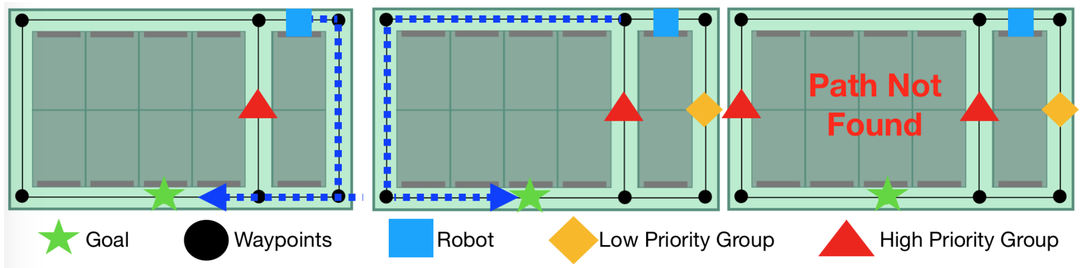

Figure 5: These are the scenarios for our case study which range from easy (one high priority group), medium (one high and low priority group), and hard (two high and one low priority group) in terms of planning difficulty (from left to right) . The dotted blue line with an arrow indicates the planned path of the robot. The algorithm finds a path in the easy and medium scenarios, but fails to find a path in the difficult scenario as it tries not to interrupt any groups.

Discussion

In this paper, we provided domain knowledge about the ED to help researchers design robots for emergence medicine. Using the insights we provided in this paper, we discussed how important it is for AI researchers to consider the severity of a patient's condition. We presented a case study of a robot that delivers materials in the ED, which uses reinforcement learning to plan paths while taking patient priority into account.

In this new, exciting domain of research, there are many opportunities to build robots to improve patient outcomes in the ED. In the future, we plan to continue our research on designing robots for the ED by developing new techniques to plan within this complex environment. For instance, we provide a simple example of a possible approach in this paper, but EDs are much larger and complex, have hundreds of patients, ED staff, and family members, and are dynamic. Recent approaches, like deep RL, might be well-suited to handle such high-dimensional state spaces, and we look forward to exploring them further.

Moving forward, we plan to conduct realistic experiments in a medical simulation and training center, which will enable us to physically simulate actual ED complexity while providing a rich testbed for our robots. For instance, we could simulate a patient in cardiac arrest, and the robot could deliver the necessary supplies to the clinicians while they begin CPR.

We hope that our work inspires robotics researchers to get involved in this domain as there is a great need to address problems in emergency medicine and explore novel technologies that can potentially save lives.

References

[1] Abdulwahid, M. A.; Booth, A.; Turner, J.; and Mason, S. M. 2018. Understanding better how emergency doctors work. analysis of distribution of time and activities of emergency doctors: a systematic review and critical appraisal of time and motion studies. *Emerg Med J* 35(11):692–700.

[2] Ahn, H. S.; Lee, M. H.; and MacDonald, B. A. 2015. Healthcare robot systems for a hospital environment: Carebot and receptionbot. In *2015 24th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, 571–576. IEEE.

[3] Alemdar, H., and Ersoy, C. 2010. Wireless sensor networks for healthcare: A survey. *Computer networks* 54(15):2688–2710.

[4] Aranibar, D. B., and Alsina, P. J. 2004. Reinforcement learning-based path planning for autonomous robots. In *EnRI-XXIV Congresso da Sociedade Brasileira de Computaç ao*, 10.

[5] Asadi-Aghbolaghi, M.; Clapes, A.; Bellantonio, M.; Escalante, H. J.; Ponce-López, V.; Baró, X.; Guyon, I.; Kasaei, S.; and Escalera, S. 2017. A survey on deep learning based approaches for action and gesture recognition in image sequences. In *2017 12th IEEE international conference on automatic face & gesture recognition (FG 2017)*, 476–483. IEEE.

[6] Aulinas, J.; Petillot, Y. R.; Salvi, J.; and Lladó, X. 2008. The slam problem: a survey. *CCIA* 184(1):363–371.

[7] Avci, A.; Bosch, S.; Marin-Perianu, M.; Marin-Perianu, R.; and Havinga, P. 2010. Activity recognition using inertial sensing for healthcare, wellbeing and sports applications: A survey. In *23th International conference on architecture of computing systems 2010*, 1–10. VDE.

[8] Bae, H.; Kim, G.; Kim, J.; Qian, D.; and Lee, S. 2019. Multi-robot path planning method using reinforcement learning. *Applied Sciences* 9(15):3057.

[9] Bragard, I.; Dupuis, G.; and Fleet, R. 2015. Quality of work life, burnout, and stress in emergency department physicians: a qualitative review. *European Journal of Emergency Medicine* 22(4):227–234.

[10] Brown, D. S., and Niekum, S. 2019. Machine teaching for inverse reinforcement learning: Algorithms and applications. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 7749–7758.

[11] Brown, D. S.; Goo, W.; Nagarajan, P.; and Niekum, S. 2019. Extrapolating beyond suboptimal demonstrations via inverse reinforcement learning from observations. *arXiv preprint arXiv:1904.06387*.

[12] Brys, T.; Harutyunyan, A.; Suay, H. B.; Chernova, S.; Taylor, M. E.; and Nowé, A. 2015. Reinforcement learning from demonstration through shaping. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.

[13] Cabibihan, J.; Javed, H.; Ang, M.; and Aljunied, S. M. 2013. Why robots? a survey on the roles and benefits of social robots in the therapy of children with autism. *International journal of social robotics* 5(4):593–618.

[14] Center, S. M. N. 2019. Dementia beyond drugs: Changing the culture of care.

[15] Chen, T. L., and Kemp, C. C. 2010. Lead me by the hand: Evaluation of a direct physical interface for nursing assistant robots. In *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 367–374. IEEE.

[16] Chisholm, C. D.; Weaver, C. S.; Whenmouth, L.; and Giles, B. 2011. A task analysis of emergency physician activities in academic and community settings. *Annals of emergency medicine* 58(2):117–122.

[17] Cobo, L. C.; Subramanian, K.; Isbell Jr, C. L.; Lanterman, A. D.; and Thomaz, A. L. 2014. Abstraction from demonstration for efficient reinforcement learning in high-dimensional domains. *Artificial Intelligence* 216:103–128.

[18] Fernando, J. B.; Tanigawa, T.; Naito, E.; Yamagami, K.; and Ozawa, J. 2012. Collision avoidance path planning for hospital robot with consideration of disabled person's movement characteristic. In *2012 IEEE 1st Global Conference on Consumer Electronics (GCCE)*, 392–396. IEEE.

[19] Frank, A., and Riek, L. 2019. Socially intelligent task and motion planning for human-robot interaction. In *In Proceedings of Robotics, Science, and Systems (RSS) Workshop on Robust Task and Motion Planning.*

[20] Franklin, A.; Liu, Y.; Li, Z.; Nguyen, V.; Johnson, T. R.; Robinson, D.; Okafor, N.; King, B.; Patel, V. L.; and Zhang, J. 2011. Opportunistic decision making and complexity in emergency care. *Journal of biomedical informatics* 44(3):469–476.

[21] Gonzales, M. J.; Cheung, V. C.; and Riek, L. D. 2015. Designing collaborative healthcare technology for the acute care workflow. In *Proceedings of the 9th International Conference on Pervasive Computing Technologies for Healthcare*, 145–152. ICST (Institute for Computer Sciences, Social-Informatics and ....

[22] Gonzales, M. J. 2016. *Safety-critical healthcare technology design*. University of Notre Dame.

[23] Griffith, S.; Subramanian, K.; Scholz, J.; Isbell, C. L.; and Thomaz, A. L. 2013. Policy shaping: Integrating human feedback with reinforcement learning. In *Advances in neural information processing systems*, 2625–2633.

[24] Hollingsworth, J. C.; Chisholm, C. D.; Giles, B. K.; Cordell, W. H.; and Nelson, D. R. 1998. How do physicians and nurses spend their time in the emergency department? *Annals of emergency medicine* 31(1):87–91.

[25] Hu, Y.; Kim, E. G.; Cao, G.; Liu, S.; and Xu, Y. 2014. Physiological acoustic sensing based on accelerometers: A survey for mobile healthcare. *Annals of biomedical engineering* 42(11):2264–2277.

[26] Jeong, S.; Logan, D. E.; Goodwin, M. S.; Graca, S.; O'Connell, B.; Goodenough, H.; Anderson, L.; Stenquist, N.; Fitzpatrick, K.; Zisook, M.; et al. 2015. A social robot to mitigate stress, anxiety, and pain in hospital pediatric care. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction Extended Abstracts*, 103–104. ACM.

[27] Jurgenson, T., and Tamar, A. 2019. Harnessing reinforcement learning for neural motion planning. *arXiv preprint arXiv:1906.00214*.

[28] Kee, R.; Knott, J. C.; Dreyfus, S.; Lederman, R.; Milton, S.; and Joe, K. 2012. One hundred tasks an hour: an observational study of emergency department consultant activities. *Emergency Medicine Australasia* 24(3):294–302.

[29] Kesiilwe, A. B.; Lederman, R.; Dreyfus, S.; and Knott, J. 2018. A time motion study evaluating the impact of information technology on emergency department consultants activities. In *ECIS*, 29.

[30] Khan, A.; Tolstaya, E.; Ribeiro, A.; and Kumar, V. 2019. Graph policy gradients for large scale robot control. *Conference on Robot Learning (CoRL).*

[31] Kolobov, A.; Peres, Y.; Lu, C.; and Horvitz, E. J. 2019. Staying up to date with online content changes using reinforcement learning for scheduling. In *Advances in Neural Information Processing Systems*, 579–589.

[32] Kubota, A.; Iqbal, T.; Shah, J. A.; and Riek, L. D. 2019. Activity recognition in manufacturing: The roles of motion capture and sEMG+inertial wearables in detecting fine vs. gross motion. In *International Conference on Robotics and Automation*. IEEE.

[33] Lara, O. D., and Labrador, M. A. 2012. A survey on human activity recognition using wearable sensors. *IEEE communications surveys & tutorials* 15(3):1192–1209.

[34] Lee, H. S., and Kim, J. 2018. Scenario-based assessment of user needs for point-of-care robots. *Healthcare informatics research* 24(1):12–21.

[35] Lee, M. H.; Ahn, H. S.; Wang, K.; and MacDonald, B. A. 2014. Design of a healthcare sensor managing system for vital sign measuring devices. In *International Conference on Simulation, Modeling, and Programming for Autonomous Robots*, 521–532. Springer.

[36] Lei, X.; Zhang, Z.; and Dong, P. 2018. Dynamic path planning of unknown environment based on deep reinforcement learning. *Journal of Robotics* 2018.

[37] Mache, S.; Vitzthum, K.; Klapp, B. F.; and Groneberg, D. A. 2012. Doctors' working conditions in emergency care units in germany: a real-time assessment. *Emerg Med J* 29(12):e1–e1.

[38] Matsumoto, S., and Riek, L. 2019. Fluent coordination in proximate human robot teaming. In *In Proceedings of the Robotics, Science, and Systems (RSS) Workshop on AI and Its Alternatives for Shared Autonomy in Assistive and Collaborative Robotics*.

[39] Mead, R., and Matarić, M. J. 2017. Autonomous human–robot proxemics: socially aware navigation based on interaction potential. *Autonomous Robots* 41(5):1189–1201.

[40] Mukai, T.; Onishi, M.; Odashima, T.; Hirano, S.; and Luo, Z. 2008. Development of the tactile sensor system of a human-interactive robot "ri-man". *IEEE Transactions on Robotics* 24(2):505–512.

[41] Narayanan, V. K.; Spalanzani, A.; Pasteau, F.; and Babel, M. 2015. On equitably approaching and joining a group of interacting humans. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 4071–4077. IEEE.

[42] Nigam, A., and Riek, L. D. 2015. Social context perception for mobile robots. In *2015 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, 3621–3627. IEEE.

[43] Onofri, L.; Soda, P.; Pechenizkiy, M.; and Iannello, G. 2016. A survey on using domain and contextual knowledge for human activity recognition in video streams. *Expert Systems with Applications* 63:97–111.

[44] Panov, A. I.; Yakovlev, K. S.; and Suvorov, R. 2018. Grid path planning with deep reinforcement learning: Preliminary results. *Procedia computer science* 123:347–353.

[45] Penaloza, C. I.; Chernova, S.; Mae, Y.; and Arai, T. 2013. Robot reinforcement learning using crowdsourced rewards. In *Proceedings of the IROS Workshop on Cloud Robotics*. Citeseer.

[46] Qureshi, A. H.; Nakamura, Y.; Yoshikawa, Y.; and Ishiguro, H. 2016. Robot gains social intelligence through multimodal deep reinforcement learning. In *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*, 745–751. IEEE.

[47] Riek, L. D. 2016. Robotics technology in mental health care. In *Artificial intelligence in behavioral and mental health care*. Elsevier. 185–203.

[48] Riek, L. D. 2017. Healthcare robotics. *Communications of the ACM* 60(11):68–78.

[49] Ronao, C. A., and Cho, S. 2016. Human activity recognition with smartphone sensors using deep learning neural networks. *Expert systems with applications* 59:235–244.

[50] Saraydaryan, J.; Jumel, F.; and Simonin, O. 2015. Robots delivering services to moving people: Individual vs. group patrolling strategies. In *2015 IEEE International Workshop on Advanced Robotics and its Social Impacts (ARSO)*, 1–8. IEEE.

[51] Subetha, T., and Chitrakala, S. 2016. A survey on human activity recognition from videos. In *2016 International Conference on Information Communication and Embedded Systems (ICICES)*, 1–7. IEEE.

[52] Takayama, L.; Dooley, D.; and Ju, W. 2011. Expressing thought: improving robot readability with animation principles. In *2011 6th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 69–76. IEEE.

[53] Tang, Y. C.; Zhang, J.; and Salakhutdinov, R. 2019. Worst cases policy gradients. *Conference on Robot Learning (CoRL)*.

[54] Tao, C.; Okafor, N.; Mehta, A.; Maddow, C.; Robinson, D.; King, B.; Zhang, J.; and Franklin, A. 2014. A work domain ontology for modeling emergency department workflow. In *7th International Workshop on Semantic Web Applications and Tools for Life Sciences, SWAT4LS 2014*. CEUR-WS.

[55] Taylor, A.; Lee, H. R.; Kubota, A.; and Riek, L. D. 2019. Coordinating clinical teams: Using robots to empower nurses to stop the line. *Proceedings of the ACM on Human-Computer Interaction* 3(CSCW):221.

[56] Taylor, A.; Chan, D.; ; and Riek, L. 2020. Robot-centric perception of human groups. *Transactions on Human Robot Interaction*.

[57] Unhelkar, V. V.; Lasota, P. A.; Tyroller, Q.; Buhai, R.; Marceau, L.; Deml, B.; and Shah, J. A. 2018. Human-aware robotic assistant for collaborative assembly: Integrating human motion prediction with planning in time. *IEEE Robotics and Automation Letters* 3(3):2394–2401.

[58] Vartiainen, E.; Husøy, K.; and Heyer, C. 2013. Better safe than sorry: collaboration in safety-critical environments. In *Proceedings of the 2013 conference on Computer supported cooperative work companion*, 327–332. ACM.

[59] Vroon, J.; Joosse, M.; Lohse, M.; Kolkmeier, J.; Kim, J.; Truong, K.; Englebienne, G.; Heylen, D.; and Evers, V. 2015. Dynamics of social positioning patterns in group-robot interactions. In *2015 24th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, 394–399. IEEE.

[60] Wada, K.; Shibata, T.; Musha, T.; and Kimura, S. 2008. Robot therapy for elders affected by dementia. *IEEE Engineering in medicine and biology magazine* 27(4).

[61] Wang, V. H., and Osborne, T. F. 2017. Social robots and other relational agents to improve patient care. *Using Technology to Improve Care of Older Adults* 166.

[62] Watkins, C.; H., J. C.; and Dayan, P. 1992. Q-learning. *Machine learning* 8(3-4):279–292.

[63] Westbrook, J. I.; Li, L.; Hooper, T. D.; Raban, M. Z.; Middleton, S.; and Lehnbom, E. C. 2017. Effectiveness of a 'do not interrupt'bundled intervention to reduce interruptions during medication administration: a cluster randomised controlled feasibility study. *BMJ Qual Saf* 26(9):734–742.

[64] Xenex. 2019. Xenex disinfection systems.

[65] Zeger, S. L.; Irizarry, R.; and Peng, R. 2006. On time series analysis of public health and biomedical data. *Annu. Rev. Public Health* 27:57–79.

[66] Zhang, S.; Wei, Z.; Nie, J.; Huang, L.; Wang, S.; and Li, Z. 2017. A review on human activity recognition using vision-based method. *Journal of healthcare engineering* 2017.