

Advancing Socially-Aware Navigation for Public Spaces

Roya Salek Shahrezaie¹ Bethany N. Manalo² Aaron G. Brantley¹ Casey R. Lynch² David Feil-Seifer¹

Abstract—Mobile robots must navigate efficiently, reliably, and appropriately around people when acting in shared social environments. For robots to be accepted in such environments, we explore robot navigation for the social contexts of each setting. Navigating through dynamic environments solely considering a collision-free path has long been solved. In human-robot environments, the challenge is no longer about efficiently navigating from one point to another. Autonomously detecting the context and adapting to an appropriate social navigation strategy is vital for social robots' long-term applicability in dense human environments. As complex social environments, museums are suitable for studying such behavior as they have many different navigation contexts in a small space.

Our prior Socially-Aware Navigation model considered context classification, object detection, and pre-defined rules to define navigation behavior in more specific contexts, such as a hallway or queue. This work uses environmental context, object information, and more realistic interaction rules for complex social spaces. In the first part of the project, we convert real-world interactions into algorithmic rules for use in a robot's navigation system. Moreover, we use context recognition, object detection, and scene data for context-appropriate rule selection. We introduce our methodology of studying social behaviors in complex contexts, different analyses of our text corpus for museums, and the presentation of extracted social norms. Finally, we demonstrate applying some of the rules in scenarios in the simulation environment.

I. INTRODUCTION

As robots continue performing a variety of non-social functions in public spaces still occupied by humans (e.g., factory, warehouse, airport), opportunities exist for robots to significantly improve the quality of our lives and contribute positively to the safety, creative potential, and atmosphere of public social spaces [1]. In studying the use of social robotics in public spaces for role-specific functions (such as tour guiding) we investigate the navigational behaviors of those robots appropriate to designated spaces. Our recent Socially-Aware Navigation (SAN) approach works to jointly optimize for navigational performance and social performance—an approach which uses a non-linear optimization over a set of objectives [2]. Primary navigation objectives may include “minimize the time to goal,” as well as social objectives, like “don't walk too closely to other people.” We have since augmented that approach by selecting these objectives based on detected environmental features [3], resulting in adaptable

navigation behavior for the space presented (i.e., hallway, art gallery, queue). Lastly, we added an enhancement that utilizes a navigation language knowledge graph to select objectives based on detected objects and people and their relationship to navigation rules [4].

Inverse Reinforcement Learning (IRL)-based approaches learn the agent's policy from demonstrated behavior [5], [6], [7]. They created a machine-learning framework that follows the behavior of a human in some tasks and learns the goal that the human is trying to reach. However, the main problem when converting a complex task into a simple reward function is that “a provided policy may be optimal for many different reward functions. Even though we have the actions from an expert, there are many different reward functions that the specialist might be attempting to maximize” [8].

These approaches are promising in a single context and can be trained to operate on multi-context designs but require a lot of human training data. The work on multi-context socially-aware navigation [9] generates social trajectories for an autonomously sensed context; it is constrained and does not maintain a knowledge base that can scale.

Socially-Aware Navigation should acquire information from the environment to detect the context in order to select environmentally-appropriate behavior. Context detection has mainly relied on deep learning techniques such as convolutional neural networks [10], which has been successful in object detection, people detection, pose detection, etc. However, in semantic navigation, knowledge is stored in the relation of concepts such as objects, utilities, or space type, then uses the relationships between these concepts to understand the context better.

Prior work relied on a set of defined social norms used as rules for navigation. These norms originated from general knowledge, demonstrating the need for continued research on identifying human-human, social interactions in various public spaces and on understanding how human-robot, social interactions may impact each space. Following the example of a museum tour, social norms were created as guidelines for navigating a more complex social environment by using identifiable tour practices. These practices were observed from real-world experiences and include, but are not limited to, acknowledging tour participants, leading participants through crowds, and being socially aware of one's surroundings [11], [12]. Additionally, robot navigation behaviors should follow social cues to identify potentially collaborative interactions between museum employees and robots.

There has been no direct and easy-to-find reference for finding social policies and norms of specific public spaces. Researchers have reverted to using social rules limited by

¹Socially Assistive Robotics Group, Department of Computer Science and Engineering, University of Nevada, Reno. 1664 N. Virginia Street, Reno, NV 89557-0171, USA rsalek@nevada.unr.edu, aaronbrantley@nevada.unr.edu, dave@cse.unr.edu

²Department of Geography, University of Nevada, Reno. 1664 N. Virginia Street, Reno, NV 89557-0171, USA bmanalo@unr.edu, caseylynch@unr.edu

general knowledge of said spaces, and in the past, others have turned to Google search results. In many cases, these rules did not match real-world etiquette or the end user's needs. For instance, when different parties in an organization have different aims, their perceptions of the robot change accordingly [13].

This paper presents a novel approach by adding an ontology to our SAN system (see Fig 1). Since social rules play an essential role in our method, we also show the pipeline for investigating socially appropriate rules for navigating in public social environments that are usually crowded and require more research on understanding their norms. We can learn important navigation information by gathering data through unstructured interviews and analyzing that language corpus. The interviews conducted for this project aimed to understand socially appropriate and inappropriate paths, movements, and behaviors within particular spaces, such as a museum. Interviews also identified how social robots should enhance the user experience once properly integrated. This framework is applicable to various contexts of different public spaces since general rules of a context are defined through Natural Language Processing (NLP) on text data from real-world scenarios and experiences. The rest of the paper is organized as related work, methods, results, implementation and usage, and discussion and future work.

II. RELATED WORK

Robots designed to share a workspace with people need to consider human comfort and safety for their long-term acceptance in public places (e.g., factories, hospitals, museums). Traditional mobile robot navigation aims to find the shortest path to its goal without considering if such a performance-oriented path is optimal for social objectives like human comfort and safety. Recently, researchers possessed social costs in mobile robot path planning to keep proper interaction distance [14], avoiding personal space [15], avoiding passing behind a person [16], avoiding activity spaces, and waiting in a line to approach a human, and having a preferred passing side [10].

Knowledge-based methods have applications for Socially-Aware Navigation. Semantic awareness has presented new opportunities for robot navigation, enabling more powerful generalization tools in representing information [17]. A location-based mobile service was developed and evaluated as an indoor navigation service [18], which helped people navigate physical difficulties. This work uses navigation context to enhance navigation behavior similar to ours, but our application is autonomous robot navigation instead of an online service for people with disabilities. In a similar work [19], the authors propose a knowledge engine that learns and shares knowledge representations for robots to complete various responsibilities.

Robots and other advanced systems were only lightly considered for museums prior to the COVID-19 pandemic; either to display in technology-driven exhibits or to provide directions to visitors as needed [20]. Directors of both the *W. M. Keck Earth Science and Mineral Engineering Museum*

and the *John and Geraldine Lilley Museum of Art* on the University of Nevada, Reno requested the development of a robot docent (industry preferred term: *tour assistant* or *robot tour guide*) for their facilities as social distancing guidelines limited their operations. Generally, a museum docent is responsible for conducting tours and providing an educational and inspirational experience to visitors. Docents complete extensive training and are strictly volunteer-based. Tours are often hosted by docents or gallery educators. During the COVID-19 pandemic, most docents were considered to be high-risk citizens (i.e., older adults or individuals with medical conditions) as one international poll determined that a majority of voluntary guides qualified as seniors [21], [22].

In another instance, institutions may not have accommodations for some physical or cognitive abilities, or for language or generational barriers; there is also an increasing need for providing inclusive museum experiences to all visitor types. Therefore, museums professionals may consider robots to address these limitations [11]. Neither museums, nor the researchers involved a wish to replace human docents, volunteers, or staff with robots. Yet, there is a shared interest in using robots for a variety of learning experiences [11]. Effective integration of social robotics into human-centered social spaces needs to be more carefully considered. Therefore, we anatomised museum operations and social interactions among a diverse group of stakeholders (e.g. docents, visitors, staff, etc.) to identify social norms in their respective environments. By allowing the stakeholders to participate in the study, they have contributed to the elucidation of social rules and guidelines for their spaces.

III. METHODS

This section introduces our methodology in two parts: sense-making of the museum environment and mining for related social rules; and leveraging said rules for a Socially-Aware Navigation architecture.

A. Museum Interviews

There is a gap between real-world social expectations and robotics research assumptions for SAN. To bridge that gap, we conducted a series of unstructured interviews to build a text corpus of museum environments and other social structures which provide accurate representations of human-human social encounters. The experiences and discussions shared by participants provide language-based data recorded in cleaned transcripts. With qualitative, quantitative, and language-processing analyses, social rules are established and written based on the interview corpus.

Interview participants were recruited via convenience and snowball sampling. Physical and digital flyers were shared around the University of Nevada, Reno campus as well as through online platforms such as LinkedIn, Instagram, and email. Early interview participants encouraged their friends to also participate in the study. The interviews have been conducted both remotely and in-person over the course of 15 months with participants ranging in age, gender identity, ethnicity, residing location, and number of languages spoken.

translation of robot behaviors from simulation to the real world. The specific robot that was simulated for this research was the Pioneer 3, which moves around using the ROS navigation stack. Since leg detection is done through finding cylinders, people were represented by cylindrical objects.

Our model allows the implementation of socially-aware behavior using the ROS navigation stack. For example, context detection output is “gallery” and rule selector has executed the *engaging* navigation behavior, the system is configured to follow that rule. This is done by getting data of people’s locations through leg detection, having the robot make decisions based on the selected rule, and finally controlling the robot’s speed based on the rule’s specifications. This system provides a base for socially-aware behavior in robots.

IV. RESULTS

This section presents three reports for analyzing methods from the interview corpus; qualitative analysis, quantitative analysis, and the natural language processing approach.




Fig. 2: Social behaviors were identified using two axes and interview text data with social rules for each behavior.

A. Identifying Activity Space Vulnerability

“Activity Space” was defined as the space between a work of art (or other artifact) and a person’s line of sight. When asked to describe their most recent or memorable museum experiences, participants were able to reference the use of activity space during HHI [12]. One participant, an artist, described how their work was intended to be touched and interacted with in a physical manner reflecting a minimally vulnerable activity space where as other participants described their museum experiences as “sterile” or uninviting indicating a more delicate or vulnerable activity space. The connections participants have between them and the artwork is rooted in individuality and artist intention.

This required us to build a spectrum of vulnerability that can be applied to a set activity space in the scene. For instance, the museum environment can be designed to engage the audience in a physical matter or to educate using visual cues or storytelling. High levels of physical engagement were linked with low vulnerability of the activity space which was confirmed by four other participants that work for museums.

Museums can have fragile atmospheres in which sounds, mechanical issues, and talking can easily disrupt someone’s activity space. Participants reported that sound was one of the most distracting elements of their museum experience mentioning, “Sometimes they play music, which is not appropriate...it can add to a piece and you don’t want that” [sic]. When asked if “art museums and galleries are quiet and easy to navigate,” participants agreed. Because more quiet activity spaces are prone to disruption by sudden movements or sounds, they can be rated with high vulnerability on the Social Behaviors Graph (Fig. 2).

B. Recognizing Density of Humans in the Scene

Participants shared their experiences interacting with humans in the museum environment [12]. Some stated other visitors were equally a part of the museum experience as the art clarifying, “I see myself and other viewers...as team members” and “I feel like museums are public spaces...there’s going to be other people.” While it is understood that other people can be a part of the experience, the density (number of people) in the scene directly affect the appropriateness of certain social movements and behaviors. Personal space and social awareness were repeatedly mentioned by participants providing insight to how visitors move through museums and prefer to interact with others. For the museum environment, participants reported high annoyance levels given scenarios involving other humans standing next to them, behind them or infiltrating their personal space without acknowledging them [12]. This was with the understanding that the space would be an open, uncrowded gallery versus a busy or tight hallway.

C. Defining the Social Behaviors

1) *Engaging*: Participant reactions to images of various humanoid robots elicited positive remarks like “cute,” “harmless,” and “delicate” [12]. If a robot fitting these remarks were integrated into an interactive exhibit with low density of people and low vulnerability in the activity space, then the robot could employ “Engaging” behaviors. Social rules for such behaviors are written as “approach disengaged visitors and ask if they have questions,” or “offer artist information to visitors in activity space.” Risk for undue harm to the activity space is low in this environment; therefore, socially appropriate movements allow for the robot to move at the speed of traffic. The robot can operate without restrictions in this mode depending on their assigned role.

2) *Conservative*: If a robot were to move slower than the speed of traffic, it can be due to a high density of people in the detected scene. Even if the activity space maintains low vulnerability, the robot should operate based

TABLE I: A list of discovered most relevant phrases used by the participants about essential concepts in our research. These are examples of social rules we found during the process.

Base Phrase	Similar Pairs	Similarity Strength
“Human-Robot Interaction	“It gonna be an interesting encounter, you know, like, a robot giving you a tour”[sic]	0.63
	“I almost feel as though with a robot, I would rather be the one to initiate conversation”[sic]	0.60
	“Using the robot for way-finding or museum information could be very beneficial in providing”[sic] guidance from point A to point B”[sic]	0.56
“Respectful Distance	Standing back farther from the image, what do we see from here or what do we see up close?”[sic]	0.49
	“My personal space most of the time is kind of screwed by the security, not the people around me”[sic]	0.46
	“And because it is an art space, it kind of like etiquette or like a norm not to stand in front”[sic]	0.43
“Appropriate Distance from a Person	So it important not to only lead the group but understand where sensitive areas to lead them to are”[sic]	0.61
	“There are people that simply don’t feel comfortable engaging and prefer to observe from the side”[sic]	0.47
	“Next to, not in front of or behind”[sic]	0.44

on the people present. This is categorized as “Conservative” given the space may be too crowded to operate without restrictions. Conservative movements may follow social rules such as “move slowly to avoid collisions,” but it would also be considered socially appropriate for a robot to offer information to interested visitors.

3) *Reserved*: Participants reported high levels of annoyance when viewing an artwork and having someone initiate a conversation [12]. This consistently spurred emotions of feeling trapped with another person as if one could not “escape” the situation when social cues are provided. When asked about when visitors would not want to be approached by a robot, participants described situations when they were alone in a space, increasing the level of vulnerability and appearance of being cornered by the robot. “Reserved” behaviors designed for SAN may reflect social rules such as “move at a consistent speed to destination” and “do not approach lone visitor” which may cause them concern. Alternate rule for social behavior quadrant may be, “allow visitors to approach robot first” to initiate an interaction.

4) *Stationary*: Considering the activity space between a visitor and an artwork, the space can be identified as highly vulnerable due to the potential of annoying a visitor or damaging an artifact in a museum [12]. While some expressed annoyance for lack of social awareness, one participant expressed concern that the robot may damage the art if it gets too close to the work. With high activity space vulnerability and high density of people in the scene, the robot has limited options for an appropriate path. More “stationary” behaviors in SAN may be reflective of a museum security guard who stays in the gallery and is available to answer questions, offer directions and more. Person-centered rules are written as, “remain in designated area until fewer people are in scene”, while environment-centered rules follow “remain in designated area for exhibit.”

D. Quantitative Results

Twenty-two questions of the interview focused on rating Human-Human Interactions (HHI) and Human-Robot

Interactions (HRI) within the museum environment (see section III-A). This allowed us to compare the two groups of answers, HHI versus HRI, when we inquired about specific actions that may happen in the museum [12]. We ran an ANOVA test on these questions data and found that one specific question showed a statistically significant difference between the two group means as determined by one-way ANOVA($F(1,6) = 9.67, p = .02$). What we found was that specifically for the question, “You are viewing an artwork, and a navigating *person* tries to engage with you,” and on the HRI side, “You are viewing an artwork, and a navigating *robot* tries to engage with you.” Results of ANOVA test show that participants were most more annoyed when the question was about the robot in the same context. For instance, people will be more annoyed by the robot if it interrupts their space in front of a painting but they care less if a person does it.

Using this result we referred back to comments made by participants when discussing engaging HHI encounters versus the engaging HRI encounters [12]. We found participants making comments about how the experience would feel similar to “starting a conversation with a child, using limited words and concepts.” Or, “with a robot, I would rather be the one to initiate conversation” and “I’d be very annoyed...maybe if the human face turns away from the face of the robot, that’s when the robot could disengage.” These direct quotes further support the evidence that a person engaging in a conversation may be less annoying than a robot engaging in a conversation. If a robot does engage a human it is important to keep the encounter brief and provide tools for disengagement as one participant clarified. More participants stated that HHI encounters increase in levels of annoyance if they are in situations like “I’m trying to walk away and you’re not letting me” and “If I’m very, obviously concentrating on something...might be a little bit of a problem.”

E. Natural Language Processing

To expand our knowledge-based graph (ontology) with practical information related to museums and concepts in

the space, such as artwork, we used NLP analysis to extract information regarding navigation rules and norms in our specific space study. Since the size of our gathered corpus is relatively small at this time, we ran sentiment analysis and paraphrase similarity detection. As the corpus expands, we will add more in-depth analysis.

1) *Sentiment Analysis*: One crucial factor in human-robot interaction is the level of acceptance of the robot by people in the environment. Although not directly related to navigation, we were interested in finding out how participants felt about having robots working as a new technology in museums. There were several direct and indirect questions regarding introducing a robot to the museum environment, questions such as *What concerns would you have in introducing a robot into your museum environment?* or *What are some limitations you see in your institution that you may want developers to consider building into a robot?* [12] When our participants answer scale rate questions, they were prompted to explain their reasoning further, so we used that information to investigate the sentiment of each answer; negative or positive feedback can give us ideas about future performance.

We used a pre-trained transformer [27] to analyze answers. From 15 answers directly related to the idea of having a robot in a museum environment, eleven came out as positive and four as negative feedback. The sentiment score overall shows that more than 70% of feedback was more welcoming towards adding a robot in museums. We randomly selected some participant answers and found they matched the analysis; for positive, we found: *"I have no concerns about robots in museums. And I don't believe it represents a limitation at all. I think for me, I view them as a plus. An additional option and an additional feature."* On the other hand, for negatives, *"I suppose that they (robots) take a little bit less presence because often they're not as tall as humans, maybe. That would be a little creepy."*

2) *Rule Extraction Based on Phrase Similarity*: Extracting information from an unstructured corpus can be challenging. One of the ways we can test our assumption toward a specific topic in the corpus is to find sentences containing specific meanings. Sentence Transformers are used for state-of-the-art sentence embedding that can be used in finding sentences with similar meanings. We used *paraphrase-TinyBERT-L6-v2* [18] to automatically search for the phrase and sentences close to chosen phrases; for example, we are interested in finding participants' opinions related to the concept, "Human and robot interaction," this sentence is input, and the model finds sentences about the same concept. We present three base phrases and extracted sentences in (Table I) that their similarity scores was higher than (0.40).

We found consistency between similar sets of each group. These results made it easier to define rules based on the interviews. However, we have not used these set of rules directly in our implementation at this point since our corpus is expanding as we conduct more interviews. Future work will present the usage of these results in our ontology and implementation.




Fig. 3: **Engaging**: The robot is too close to the person at the top right (it was previously interacting with them but further interaction could cause discomfort), so find another person to interact with. Here, the robot chooses to interact with the disengaged person at the left since the detection of their legs has the highest reliability (chance that a person is actually there). The robot begins interaction by setting a navigation goal halfway between itself and this person and moving towards it at a fast pace to match the social setting.

V. IMPLEMENTATION AND USAGE

This section presents a part of our implementation. We tested four scenarios based on our qualitative results (see section IV-C) in simulation. We will use NLP extracted rules in subsequent steps and test more scenarios. The robot uses the information from its surroundings, such as the number of people detected through laser scans and level of vulnerability of the displays in the museum space specified manually, to decide which behavior is most appropriate and act accordingly.

Engaging: Interactive exhibits that are not densely populated could be complemented by a robot using engaging navigation behavior. The fast-paced, person-to-person movement requires a casual social setting where museum attendees would not get distracted by the robot. During this behavior, the robot uses laser scans to get the locations of each person in the museum space. The robot then attempts to set a navigation goal halfway between a person and the robot in descending order of leg detection reliability. Suppose the robot fails to find a person to fulfill the interactive aspect of engaging behavior with. In that case, the robot "wanders" to a nearby random location. "Wandering" is accomplished by setting a navigation goal no more than 2 meters from the robot's current location. This recovery behavior contributes to the robot's "cute" and "harmless" presentation. If the robot's current position is too close to a person, it will not move any closer to them to avoid causing discomfort. This logic also keeps every person in the museum space engaged, as the robot will approach "disengaged" people after each interaction. The robot moves at its fastest speed during this behavior to conform to the interactive nature of the museum space (see Fig 3).




Fig. 4: **Conservative:** The robot sets a goal halfway between the person at the top right and itself, reflecting the interactive but slightly crowded social setting. The robot chooses not to move towards the people at the bottom left and right because they are considered “too close” to the robot and interaction with them could cause discomfort due to inappropriate navigation behavior. The robot moves at a medium-fast pace towards this destination to demonstrate the expected etiquette of a densely populated museum space with low vulnerability.

Conservative: Conservative behavior should occur when a museum space is densely populated but has a low vulnerability. This is comparable to engaging behavior as the robot seeks to interact with others in the museum to reflect the low vulnerability of the space. However, the robot stays further away from other people in the museum and moves slightly slower to avoid collisions. If there are no people to interact with, the robot will stay in place instead of “wandering.” This reflects the behavioral expectations of the museum setting, where people might interact but will not interfere with each other. Like engaging behavior, the robot will end interaction with a person if they are too close to the robot’s current location, but the area that counts as “too close” is twice as large during this behavior. This results in the robot keeping a greater distance from every person in the room. The robot moves at a medium-fast speed to match the social setting of the museum space (see Fig 4).

Reserved: The robot will act as if no people are present during reserved behavior, moving to one of the predetermined locations set in the map. This socially detached navigation method demonstrates the behavioral expectations of a sparsely populated, high vulnerability museum space, where people should not initiate interaction with others unless it is explicitly desired. Once the robot reaches this destination (or if it is already near a predetermined location), it will stay in place, not creating any new paths. This keeps the robot from approaching lone visitors and allows museum attendees that do wish to interact with the robot to do so. Similar to conservative behavior, if the robot cannot create a path to a predetermined location, it will stay in place until it can do so. The robot moves at a medium-slow pace during this behavior to make it apparent that the robot is not trying to




Fig. 5: **Reserved:** The robot finds itself in a sparsely populated setting with high vulnerability and utilizes reserved navigation behavior as a result. To prevent discomfort, the robot keeps its distance from other people in the museum by avoiding direct interaction and restricting movement to programmer-defined locations in the museum. Here, the predetermined location is at the center of the museum space, so the robot will travel to this location at a medium-slow pace and remain there, standing by for human-initiated interaction.

engage with another person in the museum (see Fig 5).

Stationary: Like reserved behavior, the robot’s navigation is not directly influenced by other people during stationary navigation behavior. This is done to match the high density and vulnerability of the museum space that warrants this behavior. Museum attendees would likely become irritated if




Fig. 6: **Stationary:** The robot travels to its single stationary location where it will assume the role of a kiosk. As it travels towards this secluded location, the robot moves at a slow pace to prevent diversion of the museum attendees’ attention. To complement the high vulnerability and density of the museum space, people will have to approach the robot in order to initiate interaction. The robot will remain at this destination until the museum space warrants the start of another behavior (either people leave the museum or the environment becomes less vulnerable).

a robot were constantly moving around the space. There is only one predetermined location during stationary behavior, which is where the robot will serve a role similar to a security guard by giving directions or answering questions. Ideally, it would be in a secluded area that would not distract those in the museum. Instead of the robot approaching people to initiate interaction, people will have to approach the robot. Like reserved behavior, the robot will remain in place until it can create a path to the predetermined location. When the robot can travel to this location, it will move there slowly to avoid diverting the museum attendees' attention from the displays to itself (see Fig 6).

VI. DISCUSSION AND FUTURE WORK

Design challenges posed by robots added in public places go beyond conventional concerns of aesthetics and usability. Relevant social rules must be defined through real-world experiences shared through interviews or other text resources. This paper presented a framework for gathering language-based information focused on real-world experiences through interviews with experienced participants from museums. We also applied language based methods for analyzing our text data. Based on our findings, we implemented a few scenarios and developed navigation for robots in public places.

Future research will include the expansion of the text corpus with the addition of more interviews with more related questions to navigation. Additionally, we aim to use knowledge extraction methods based on knowledge graphs to use the relation of environment elements in rule selection. Finally, we will add more scenarios to our experiment results in simulation and real-world implementation and evaluate resultant behavior with observers/participants using perceived social intelligence (PSI) [28].

ACKNOWLEDGEMENT

The authors would like to acknowledge the financial support of this work by the National Science Foundation (NSF #SES-2121387, #IIS-1719027, and #IIS-1757929).

REFERENCES

- [1] M. Mintrom, S. Sumartojo, D. Kulić, L. Tian, P. Carreno-Medrano, and A. Allen, "Robots in public spaces: implications for policy design," *Policy Design and Practice*, vol. 0, no. 0, pp. 1–16, 2021. [Online]. Available: <https://doi.org/10.1080/25741292.2021.1905342>
- [2] S. Forer, S. B. Banisetty, L. Yliniemi, M. Nicolescu, and D. Feil-Seifer, "Socially-aware navigation using non-linear multi-objective optimization," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1–9.
- [3] S. B. Banisetty, V. Rajamohan, F. Vega, and D. Feil-Seifer, "A deep learning approach to multi-context socially-aware navigation," in *International Conference on Robot and Human Interactive Communication (RO-MAN)*, Vancouver, BC, Canada, August 2021, pp. 23–30.
- [4] R. Salek Shahrezaie, S. B. Banisetty, M. Mohammadi, and D. Feil-Seifer, "Towards deep reasoning on social rules for socially aware navigation," in *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, 2021, pp. 515–518.
- [5] A. Y. Ng, S. J. Russell, et al., "Algorithms for inverse reinforcement learning," in *icml*, vol. 1, 2000, p. 2.
- [6] H. Kretschmar, M. Spies, C. Sprunk, and W. Burgard, "Socially compliant mobile robot navigation via inverse reinforcement learning," *The International Journal of Robotics Research*, vol. 35, no. 11, pp. 1289–1307, 2016.
- [7] R. Ramon-Vigo, N. Perez-Higueras, F. Caballero, and L. Merino, "Transferring human navigation behaviors into a robot local planner," in *Robot and Human Interactive Communication, 2014 RO-MAN: The 23rd IEEE International Symposium on*. IEEE, 2014, pp. 774–779.
- [8] J. Alexander, "learning from humans: what is inverse reinforcement learning?," *The Gradient*, 2018.
- [9] S. B. Banisetty, "Multi-context socially-aware navigation using non-linear optimization," Ph.D. dissertation, University of Nevada, Reno, 2020.
- [10] S. B. Banisetty, S. Forer, L. Yliniemi, M. Nicolescu, and D. Feil-Seifer, "Socially aware navigation: A non-linear multi-objective optimization approach," *ACM Transactions on Interactive Intelligent Systems*, vol. 11, no. 2, July 2021.
- [11] L. B. Manalo and C. R. Lynch, "Integrating social robotics into the art museum environment: Guidelines for facilitating a meaningful museum experience," 2021, unpublished paper.
- [12] L. B. Manalo, Muñoz-Viso, R. Salek Shahrezaie, C. R. Lynch, and D. Feil-Seifer, "Museum robot project interview protocol." [Online]. Available: https://docs.google.com/document/d/1Tlh_9tnIPjpJuOdOTqJZEOMWOLF6K6w_Pzrtm84SapU/edit
- [13] B. Mutlu and J. Forlizzi, "Robots in organizations: The role of work-flow, social, and environmental factors in human-robot interaction," in *2008 3rd ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2008, pp. 287–294.
- [14] R. Kirby, R. Simmons, and J. Forlizzi, "Companion: A constraint-optimizing method for person-acceptable navigation," in *RO-MAN 2009-The 18th IEEE International Symposium on Robot and Human Interactive Communication*. IEEE, 2009, pp. 607–612.
- [15] D. V. Lu, D. Hersherberger, and W. D. Smart, "Layered costmaps for context-sensitive navigation," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2014, pp. 709–715.
- [16] E. A. Sisbot, L. F. Marin-Urias, R. Alami, and T. Simeon, "A human aware mobile robot motion planner," *IEEE Transactions on Robotics*, vol. 23, no. 5, pp. 874–883, 2007.
- [17] J. Crespo, J. C. Castillo, O. M. Mozos, and R. Barber, "Semantic information for robot navigation: A survey," *Applied Sciences*, vol. 10, no. 2, p. 497, 2020.
- [18] V. Tsetos, C. Anagnostopoulos, P. Kikiras, and S. Hadjiefthymiades, "Semantically enriched navigation for indoor environments," *International Journal of Web and Grid Services*, vol. 2, no. 4, pp. 453–478, 2006.
- [19] A. Saxena, A. Jain, O. Sener, A. Jami, D. K. Misra, and H. S. Koppula, "Robobrain: Large-scale knowledge engine for robots," *arXiv preprint arXiv:1412.0691*, 2014.
- [20] L. Styx, "How are museums using artificial intelligence, and is ai the future of museums?" *MuseumNext Digital*, 09 2020. [Online]. Available: <https://www.museumnext.com/article/artificial-intelligence-and-the-future-of-museums/>
- [21] C. for Disease Control and Prevention, "People at increased risk," *COVID-19*, 04 2021. [Online]. Available: <https://www.cdc.gov/coronavirus/2019-ncov/need-extra-precautions/index.html>
- [22] M. Love, "An international overview of art docents," *The Docent Educator*, vol. 3, no. 3, pp. 8–11, Spring 1994.
- [23] M. Collan and P. Luukka, *Proceedings of NSAIS16 Workshop on Intelligent and Adaptive Systems*, 08 2016.
- [24] M. Pawlaczyk-Luszczynska, A. Dudarewicz, M. Waszkowska, and M. Sliwiska-Kowalska, "Assessment of annoyance from low frequency and broadband noise," *International journal of occupational medicine and environmental health*, vol. 16, pp. 337–43, 02 2003.
- [25] S. Macenski, F. Martin, R. White, and J. G. Clavero, "The marathon 2: A navigation system," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, oct 2020. [Online]. Available: <https://doi.org/10.1109/2FIROS45743.2020.9341207>
- [26] R. Vaughan, "Massively multi-robot simulation in stage," *Swarm intelligence*, vol. 2, no. 2, pp. 189–208, 2008.
- [27] N. Reimers and I. Gurevych, "Making monolingual sentence embeddings multilingual using knowledge distillation," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 11 2020. [Online]. Available: <https://arxiv.org/abs/2004.09813>
- [28] A. Honour, S. B. Banisetty, and D. Feil-Seifer, "Perceived social intelligence as evaluation of socially navigation," in *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, ser. HRI '21 Companion. New York, NY, USA: Association for Computing Machinery, 2021, p. 519–523. [Online]. Available: <https://doi.org/10.1145/3434074.3447226>