

Speech-Guided Sequential Planning for Autonomous Navigation using Large Language Model Meta AI 3 (Llama3)

Alkesh K. Srivastava[Ⓛ] and Philip Dames[Ⓛ]

Temple University, Philadelphia PA 19122, USA
{alkesh, pdames}@temple.edu

Abstract. In social robotics, a pivotal focus is enabling robots to engage with humans in a more natural and seamless manner. The emergence of advanced large language models (LLMs) has driven significant advancements in integrating natural language understanding capabilities into social robots. This paper presents a system for speech-guided sequential planning in pick and place tasks, which are found across a range of application areas. The proposed system uses Large Language Model Meta AI (Llama3) to interpret voice commands by extracting essential details through parsing and decoding the commands into sequential actions. These actions are sent to DRL-VO, a learning-based control policy built on the Robot Operating System (ROS) that allows a robot to autonomously navigate through social spaces with static infrastructure and crowds of people. We demonstrate the effectiveness of the system in simulation experiment using Turtlebot 2 in ROS1 and Turtlebot 3 in ROS2. We conduct hardware trials using a Clearpath Robotics Jackal UGV, highlighting its potential for real-world deployment in scenarios requiring flexible and interactive robotic behaviors.

Keywords: Human-Robot Interactions · Large Language Models · Motion Planning · Natural Language Processing.

1 Introduction

Social robotics aims to enable robots and humans to cohabitate in a natural and intuitive manner. In the field of collaborative robotics, robots are required to share workspaces with humans, making it essential for robots to understand and execute human commands effectively [3]. Robots are increasingly being used in retail, industrial settings, households, and office environments. In such settings, robots are often asked to carry out tasks involving multiple steps or to complete a series of tasks in a specific order. For instance, in a retail environment, a robot might be required to navigate the store to pick up items from different aisles and deliver them to a checkout counter. In an office, a robot might need to collect documents from multiple departments and deliver them to a central location, such as Human Resources. At home, robots could be tasked with household chores that involve sequential actions such as serving food from the kitchen and

then cleaning up after dinner. Such robotic tasks are often modeled as Vehicle Routing Problems or other optimization challenges [9]. However, the complexity of these problems increases in dynamic environments, with more assigned tasks, and a growing number of robots. Robots are also increasingly deployed for these tasks in hostile or hard-to-reach areas [1]. Solving these issues is challenging on its own, and incorporating natural language understanding for collaborative robotics adds an additional layer of difficulty.

This paper aims to enable people to sequentially task robots using natural language inputs. However, in experimental robotics, models are often designed for specific tasks with rigid input formats to ensure successful task execution. This rigidity ensures that robots can accomplish their assignments accurately but does not account for the natural variability in human communication. Humans, celebrating their individuality, interact with robots in diverse and natural ways, often not conforming to predefined input formats. Therefore, for collaborative robotics to be truly successful, robots must be able to interpret a variety of human commands and translate them into actionable plans [22]. To address this need, we will build upon recent advancements in large language models, which have shown great promise in enhancing natural language understanding.

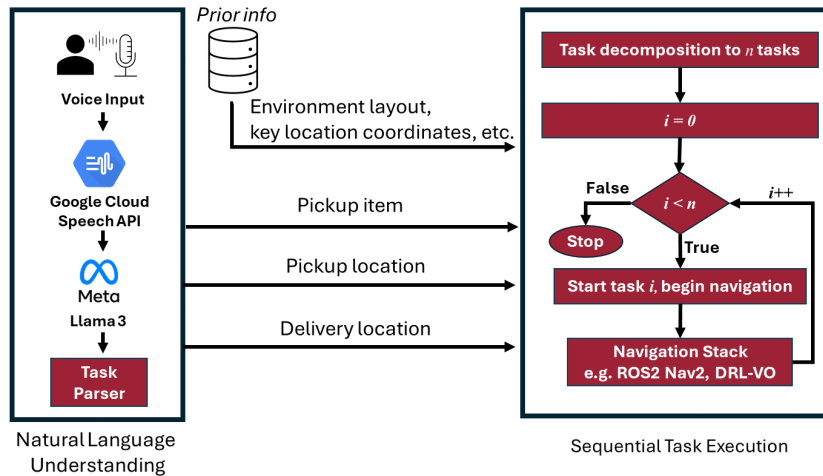


Fig. 1: Overview of the proposed system. The process begins with converting verbal commands into text using the Google Cloud Speech API. The Llama3 model processes this text to extract essential details, such as the pickup location, item, and delivery destination, using regex-based parsing. These parsed commands, along with prior environmental information, are sent to the Task Execution module, where they are translated into a sequence of actions that are then executed.

Our system, illustrated in Fig. 1, consists of two main modules. The first is the natural language understanding module, which processes user speech or text input and outputs a set of task parameters. We utilize speech-to-text tech-

nology to convert verbal commands into text. The text is fed into Llama3 [16], an autoregressive large language model, which extracts essential details—such as pickup location, item, and delivery destination—using regex-based parsing. In other words, we are using Llama3 to “translate” free-form natural language commands into a standardized, easily parsable form for robust task extraction. Although alternatives like BERT [7] could be trained on large datasets to handle this process, they would likely eliminate the need for an LLM. Instead, we prioritize Llama3’s ability to generalize and flexibly interpret commands, which allows us to maintain a simpler, rule-based parsing mechanism while still benefiting from the LLM’s advanced linguistic capabilities.

The second module translates the parsed commands into a sequence of actions by referencing a predefined dictionary of environment coordinates. This dictionary stores coordinates of key landmarks, eliminating the need for semantic planning and enabling faster execution. Task execution is managed using a finite-state machine to control the flow of operations, along with off-the-shelf robot navigation algorithms.

We demonstrate the efficacy of our system through simulated and hardware experiments using three different robot models, two navigation algorithms, and two environments. These experiments highlight the modularity of our system and the potential for practical applications in real-world scenarios.

2 Related Work

In this section, we review the literature related to each of our two modules: natural language understanding (NLU) and robotic navigation.

2.1 Natural Language Understanding

Recent advancements in large language models (LLMs) have significantly enhanced natural language understanding (NLU) capabilities across various domains, including robotics [28]. Models like GPT-3 and BERT [7] have demonstrated exceptional performance in comprehending and generating human-like text, owing to their extensive training on large datasets. By fine-tuning these models for specific tasks, researchers have achieved state-of-the-art results in tasks spanning from text classification to question answering [28,29], thereby enhancing their performance on conventional NLU benchmarks [6,10]. SayPlan [21] focuses on scalable task planning using 3D scene graphs but does not integrate speech-recognition or incorporate socially compliant algorithms such as DRLVO [27] for plan execution.

LLMs excel particularly in tasks requiring broad contextual understanding or dealing with unstructured data [6], which is pertinent to collaborative and social robotics applications. LLMs exhibit generalization capabilities, surpassing traditionally fine-tuned models in handling diverse and adversarial inputs [14,20]. In terms of practical applications, recent studies have showcased LLM effectiveness across various NLU tasks, such as machine translation [2], question answering [25], and text classification [4].

Contributions The novelty of our framework lies in its integration of speech recognition and NLU capabilities, leveraging LLMs to enhance task classification accuracy. Our approach seamlessly integrates Google Speech Cloud API for speech-to-text transcription with Llama3 for task classification.

2.2 NLU-guided Social Autonomous Navigation

The vision of seamless integration of mobile robots into human environments has been extensively studied. Robots like RHINO and MINERVA were deployed in museums and solely focused on autonomous navigation amidst humans [13]. Various approach of navigation has treated people as dynamic, non-responsive obstacles, and emphasized collision avoidance [26], while other approaches integrated human motion prediction with robot decision-making, acknowledging mutual influences between actions of robots and people [8].

Recent advancements in robotic navigation guided by natural language have explored diverse methodologies to enhance human-robot interaction and adaptability in dynamic environments. FollowNet [24] and LM-Nav [23] leverage end-to-end neural architectures and large pre-trained models to interpret natural language instructions and navigate complex environments, demonstrating success in simulated and real-world scenarios. GOAT [5] offers a multimodal navigation system that integrates language descriptions and object recognition. However, it lacks a comprehensive approach to dynamic social interactions in complex environments. Arena 3.0 [12] provides a realistic simulation environment for social navigation but does not incorporate the advanced NLU features and real-time feedback mechanisms that the proposed approach offers.

Contributions Our work advances beyond existing methodologies by combining robust NLU capabilities with adaptive navigation strategies. Unlike previous approaches, which often rely on fixed control policies or simplistic command parsing techniques, our methodology leverages the capabilities of Llama3 for robust natural language understanding and DRL-VO [27] for adaptive and socially-compliant navigation.

3 Problem Statement

The general task addressed in this work is to understand human commands using Large Language Models (LLMs) and form a sequential plan for execution. The overall problem can be subdivided into two problems:

Problem 1: NLU for Command Interpretation

This paper focuses on the scenario where a social robot receives instructions to pick up an object from a specified location and deliver it to another designated place. This task exemplifies a common class of multi-step operations in social robotics, where the robot must navigate through sequential actions such as in household assistance for cleaning tasks, delivering medication in health

care settings, etc. By addressing pickup and delivery planning, we aim to tackle fundamental challenges that underpin various practical applications.

Let C denote the natural language command (in text or speech) provided by the user, which includes information about pickup location L_{pickup} , delivery location L_{delivery} , and the pickup item I . The task of the robot is to design a Natural Language Understanding (NLU) system that accurately parses C into a structured format:

$$L_{\text{pickup}}, L_{\text{delivery}}, I = \text{NLU}(C),$$

where $\text{NLU}(C)$ process C using Llama3 and parses the obtained information to extract the task parameters L_{pickup} , L_{delivery} , and I .

Problem 2: Autonomous Navigation

We assume an environment whose layout of the static infrastructure is known and we possess coordinates for key locations within it (e.g., room numbers). The robot shares the environment with people who may move around during the robot’s operation. The robot model used in this study is a differential wheel drive robot, which is a representative model of various ground robots such as TurtleBots, Jackal UGV, and Moxi by Diligent Robotics, though the proposed system is applicable to any mobile robot.

Once the task parameters are identified by the NLU, we can look up the coordinates of the pickup and delivery locations, L_{pickup} , L_{delivery} , in the environment map. The next challenge is to navigate through the environment to reach those locations. The robot does this using a control policy π with parameters θ that selects steering actions a_t based on partial environmental observation o_t (obtained from the sensors and perception system):

$$a_t \sim \pi_{\theta}(a_t|o_t).$$

4 Methodology

In this section, we describe the methodology of the proposed Speech-Guided Sequential Planner for Autonomous Navigation. The proposed system integrates speech recognition, natural language understanding, and an advanced control technique to enable a robot to autonomously navigate in an environment populated with human pedestrians. First, we address the natural language understanding aspect of Problem 1 in Section 4.1. Followed by addressing the autonomous navigation challenge posed by Problem 2 in Section 4.2.

4.1 Natural Language Understanding (NLU)

The first module takes in a natural language input (speech or text) and parses it to extract the task information. This is done in two steps.

Speech To Text Conversion To convert speech to text, we utilize the Google Speech Recognition API [15] through the Python3 `speech_recognition` library for precise speech-to-text conversion. Initially, ambient noise levels are calibrated using the microphone’s input during the first second without processing speech, adjusting the recognizer’s sensitivity accordingly. Google’s advanced speech recognition technology employs signal processing and machine learning methods. It begins by preprocessing the audio input to extract spectral representations and phonetic patterns. Deep neural networks then map these features to textual sequences with high accuracy and efficiency.

Large Language Model (LLM) The natural language understanding (NLU) process encompasses several key steps. Firstly, the received text undergoes preprocessing, which includes tasks such as punctuation removal and text formatting to maintain consistency throughout the text. Following preprocessing, the text is fed into the Llama3 API [16] provided by Groq [11] for context understanding and semantic parsing.¹ This step involves extracting structured data essential for task execution. The extracted information is parsed using regular expressions to identify the pickup location L_{pickup} , delivery location L_{delivery} , and the item to be picked up I . We utilize the LLM as a translation layer to reformat complex, natural language commands into a standardized, structured format that simplifies parsing. The integration of Llama3 enhances the system’s capability to accurately interpret and process natural language commands, promoting seamless interaction and task execution in robotic applications. Opting for regex-based parsing over learned task classifiers ensures greater flexibility in handling diverse human commands and reduces computational costs, making it more practical for implementation in robots.

4.2 Sequential Task Assignment and Execution

Once the command is interpreted, we formulate and execute the task sequence.

Task Assignment For task assignment, our system employs a Finite State Machine (FSM) for handling simple tasks such as a single pickup and delivery. The FSM consists of states like Idle, Navigating to Pickup, Picking Up Item, Navigating to Delivery, and Delivering Item, with transitions triggered by events such as reaching a location or completing an action. This approach is efficient for straightforward tasks. However, for more complex tasks that involve multiple

¹ We employ the 8 billion parameter version (`llama3-8b-8192m`) instead of the larger counterparts, such as the 80 billion parameter model, due to larger models exhibiting a tendency to assign specific room numbers when interpreting generic commands. For instance, if asked to navigate to “TRAIL lab,” it might erroneously label it with a hallucinated room number, like “Room 111.” Additionally, we can use the smaller model onboard the robot rather than relying on remote API services, leading to quicker response times and greater practicality for real-time applications.

steps and dependencies, we can employ hierarchical task planners like GTPy-hop [17] or SHOP [18]. These planners decompose high-level tasks into manageable sub-tasks, enabling the robot to handle complex scenarios, thus justifying their integration into our system for enhanced task management and execution.

Autonomous Navigation We utilize two different navigation frameworks, depending on the situation. In static worlds (i.e., without people), we use the ROS2-Nav2 navigation stack. This is one of the most commonly used navigation frameworks in mobile robots. In environments full of people, we utilize the DRL-VO (Deep Reinforcement Learning with Velocity Obstacles) [27] navigation system, which the authors previously designed to enable robots to navigate through crowded and dynamic environments. The control policy is a convolutional neural network that uses lidar scans, maps of pedestrian locations/speeds, and goal coordinates to generate velocity commands for the robot. DRL-VO yields higher speeds and few collisions than other robot controllers (including the ROS navigation stack), especially in dense crowds.

5 Experiments

To evaluate the system, we first assess the speech-to-task classification by having volunteers self-report the accuracy of the classification on 10 natural task statements. We then evaluate the system integration with three distinct robotic experiments: a simulation in an empty office environment using a Turtlebot3 in ROS2, a simulation in a lobby with people using a Turtlebot2 in ROS1, and a hardware implementation in a lobby with people using a Clearpath Robotics Jackal UGV. The experimental setup is described in Fig. 2. These experiments collectively demonstrate the system’s versatility and effectiveness across different scenarios, from pedestrian-free areas to crowded social spaces.


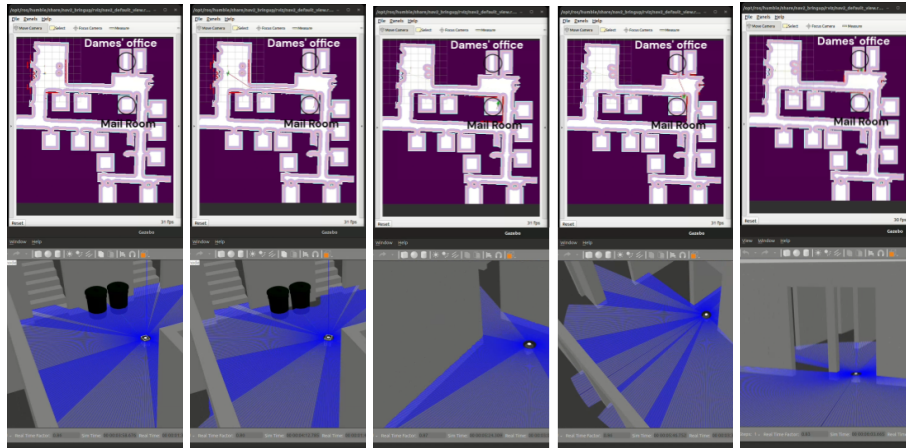


Fig. 2: (Left) The layout of the Mechanical Engineering department at Temple University, used in the ROS2-Nav2 simulation with Turtlebot3. (Center) The layout of the lobby of the College of Engineering at Temple University, which is used for both simulation experiments with Turtlebot2 using DRL-VO and hardware experiments with the Jackal UGV. (Right) The Jackal UGV.

5.1 Speech-to-Task Accuracy

To test the efficacy of the natural language understanding aspect of our system, we asked 10 volunteers to speak 10 commands and report the number of correct classifications. Out of the 10 commands, 5 were provided to them and the rest were the volunteer’s own command statements. This ensures that the accuracy reported by the system is unbiased. Examples of the commands include “Could you please bring the keys from security to TRAIL?” and “I forgot my laptop, please bring a laptop from the computer station to the robotics lab.” We observed an average accuracy of 84.37% for task classification. Feedback from the volunteers revealed that using identifiers like “the” or “a” before the location or pickup item confused the NLU system, suggesting a need for improved preprocessing to handle such variations in command phrasing.



(a) Step 1: Localizing robot’s current position in the environment (upper left room). (b) Step 2: Planning path and navigating to the pickup location L_{pickup} (Mail Room). (c) Step 3: Reaching L_{pickup} (Mail room) and executing pickup action (envelopes). (d) Step 4: Planning path and navigating to the delivery location $L_{delivery}$ (Dames’ Office). (e) Step 5: Reaching $L_{delivery}$ (Dames’ Office) and completing the sequential tasks.

Fig. 3: Simulation experiment depicting the steps of progress for pickup of item I (envelopes) from L_{pickup} (Mail Room) and its delivery to $L_{delivery}$ (Dames’ Office).

5.2 System Integration Testing

We conducted three tests, outlined in Table 1 that vary the robot platform, environment, and ROS version to demonstrate the flexibility of the system. The table also provides links to full videos of all experiments, and Fig. 3 shows a series of snapshots of the first experiment with the Turtlebot 3. The videos show that in all cases the robot is able to complete the task given to it, either from a natural language text input on a remote computer or a spoken command. The

simulated and hardware tests in the lobby environment both have pedestrians, demonstrating the system’s capability to navigate autonomously in social spaces with static infrastructure and moving crowds, which is essential for getting robots outside of controlled lab and factory environments.

Table 1: Overview of Experiments

Type	Robot	Navigation	Environment	Videos
Simulation	Turtlebot3	ROS2-Nav2	ME Dept.	https://youtu.be/L3kIdW80ZK0 https://youtu.be/HGL9E_DZUsk
Simulation	Turtlebot2	DRL-VO	Lobby	https://youtu.be/TfleIxcCoE8 https://youtu.be/BaoLZ68bkAM
Hardware	Jackal	DRL-VO	Lobby	https://youtu.be/hzYRvpX9Qe8

5.3 Discussion

Our experimental setup aimed to demonstrate the versatility and compatibility of our system with different robotic platforms and environments. In a controlled, pedestrian-free environment, we utilized the ROS2 Nav2 package with Turtlebot3, emphasizing the system’s capability to operate effectively with smaller base robots and showcasing its integration with advanced ROS2 features. In human-populated environments, where the risk of collision is higher, we used the DRL-VO [27] algorithm on a Turtlebot2 and a Jackal. This highlights our system’s adaptability to various robotic platforms, software stacks, and environments to be able to operate in a range of settings.

6 Conclusion

In this paper, we explore the feasibility and effectiveness of integrating advanced natural language processing with robotic navigation systems to enable people to use natural language to ask robots to complete multi-step tasks. We leverage Llama3 to parse commands and different off-the-shelf navigation algorithms, and we apply our system to a range of robot models, software stacks, and environments. Our research enhances social and collaborative robotics by enabling people to naturally interact with robots in shared spaces.

Future work will aim to enhance the versatility to more real-world scenarios. We will update the robotic hardware by adding a gripper to the robot. This will allow it to pick up and deliver objects as well as autonomously utilize elevators, presenting an intriguing opportunity to expand their operational capabilities in multi-story buildings and complex indoor environments. We will also update the software with the goals of: 1) enhancing the success rate of our NLU module, 2) allowing the robot to respond to the user if it does not understand the given

query [22], and 3) broadening the scope of tasks, which may include additional steps to complete, by utilizing hierarchical task networks (HTNs) [19] or similar planning paradigms.

Acknowledgment

This work was funded by NSF grant CNS-2143312. The authors would like to thank Zhanteng Xie, Jared Levin, and Alexander Derrico for their assistance with the hardware experiments.

References

1. Bellingham, J.G., Rajan, K.: Robotics in remote and hostile environments. *Science* **318**(5853), 1098–1102 (2007)
2. Brants, T., Popat, A., Xu, P., Och, F.J., Dean, J.: Large language models in machine translation. In: Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL). pp. 858–867 (2007)
3. Breazeal, C., Dautenhahn, K., Kanda, T.: Social robotics. *Springer Handbook of Robotics* pp. 1935–1972 (2016)
4. Chae, Y., Davidson, T.: Large language models for text classification: From zero-shot learning to fine-tuning. Open Science Foundation (2023)
5. Chang, M., Gervet, T., Khanna, M., Yenamandra, S., Shah, D., Min, S.Y., Shah, K., Paxton, C., Gupta, S., Batra, D., et al.: Goat: Go to any thing. arXiv preprint arXiv:2311.06430 (2023)
6. Chang, Y., Wang, X., Wang, J., Wu, Y., Yang, L., Zhu, K., Chen, H., Yi, X., Wang, C., Wang, Y., et al.: A survey on evaluation of large language models. *ACM Transactions on Intelligent Systems and Technology* **15**(3), 1–45 (2024)
7. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
8. Fridovich-Keil, D., Ratner, E., Peters, L., Dragan, A.D., Tomlin, C.J.: Efficient iterative linear-quadratic approximations for nonlinear multi-player general-sum differential games. In: 2020 IEEE international conference on robotics and automation (ICRA). pp. 1475–1481. IEEE (2020)
9. González, D., Pérez, J., Milanés, V., Nashashibi, F.: A review of motion planning techniques for automated vehicles. *IEEE Transactions on Intelligent Transportation Systems* **17**(4), 1135–1145 (2015)
10. Huang, J., Chang, K.C.C.: Towards reasoning in large language models: A survey. arXiv preprint arXiv:2212.10403 (2022)
11. Inc., G.: Groq. (<https://groq.com/>) (nd), [Online; accessed July 2024]
12. Kästner, L., Shcherbyna, V., Zeng, H., Le, T.A., Schreff, M.H.K., Osmaev, H., Tran, N.T., Diaz, D., Golebiowski, J., Soh, H., et al.: Arena 3.0: Advancing social navigation in collaborative and highly dynamic environments. arXiv preprint arXiv:2406.00837 (2024)
13. Lakemeyer, W.S.G., Schulz, D., Thrun, S.: The interactive museum tour-guide robot. In: Proc. of the 15th National Conference on Artificial Intelligence (1998)

14. Liang, P., Bommasani, R., Lee, T., Tsipras, D., Soylu, D., Yasunaga, M., Zhang, Y., Narayanan, D., Wu, Y., Kumar, A., et al.: Holistic evaluation of language models. arXiv preprint arXiv:2211.09110 (2022)
15. LLC, G.: Cloud speech-to-text API. (<https://cloud.google.com/speech-to-text>) (nd), [Online; accessed July 2024]
16. Meta AI: Meta Llama 3. (<https://github.com/meta-llama/llama3>) (2024)
17. Nau, D., Bansod, Y., Patra, S., Roberts, M., Li, R.: Gtphyop: A hierarchical goal+task planner implemented in python. HPlan 2021 p. 21 (2021)
18. Nau, D., Cao, Y., Lotem, A., Munoz-Avila, H.: The shop planning system. AI Magazine **22**(3), 91–91 (2001)
19. Nejati, N., Langley, P., Konik, T.: Learning hierarchical task networks by observation. In: Proceedings of the 23rd international conference on Machine learning. pp. 665–672 (2006)
20. Nie, Y., Williams, A., Dinan, E., Bansal, M., Weston, J., Kiela, D.: Adversarial nli: A new benchmark for natural language understanding. arXiv preprint arXiv:1910.14599 (2019)
21. Rana, K., Haviland, J., Garg, S., Abou-Chakra, J., Reid, I.D., Suenderhauf, N.: Sayplan: Grounding large language models using 3d scene graphs for scalable task planning. CoRR (2023)
22. Ren, A.Z., Dixit, A., Bodrova, A., Singh, S., Tu, S., Brown, N., Xu, P., Takayama, L., Xia, F., Varley, J., Xu, Z., Sadigh, D., Zeng, A., Majumdar, A.: Robots that ask for help: Uncertainty alignment for large language model planners. In: 7th Annual Conference on Robot Learning (2023), <https://openreview.net/forum?id=4ZK80DNyFXx>
23. Shah, D., Osinski, B., Levine, S., et al.: Lm-nav: Robotic navigation with large pre-trained models of language, vision, and action. In: Conference on robot learning. pp. 492–504. PMLR (2023)
24. Shah, P., Fiser, M., Faust, A., Kew, J.C., Hakkani-Tur, D.: Follownet: Robot navigation by following natural language directions with deep reinforcement learning. arXiv preprint arXiv:1805.06150 (2018)
25. Singhal, K., Tu, T., Gottweis, J., Sayres, R., Wulczyn, E., Hou, L., Clark, K., Pfohl, S., Cole-Lewis, H., Neal, D., et al.: Towards expert-level medical question answering with large language models. arXiv preprint arXiv:2305.09617 (2023)
26. Vasquez, D., Okal, B., Arras, K.O.: Inverse reinforcement learning algorithms and features for robot navigation in crowds: an experimental comparison. In: 2014 IEEE/RSJ International Conference on Intelligent Robots and Systems. pp. 1341–1346. IEEE (2014)
27. Xie, Z., Dames, P.: DRL-VO: Learning to navigate through crowded dynamic scenes using velocity obstacles. IEEE Transactions on Robotics **39**(4), 2700–2719 (Aug 2023). <https://doi.org/10.1109/TR0.2023.3257549>
28. Zhao, W.X., Zhou, K., Li, J., Tang, T., Wang, X., Hou, Y., Min, Y., Zhang, B., Zhang, J., Dong, Z., et al.: A survey of large language models. arXiv preprint arXiv:2303.18223 (2023)
29. Zhou, C., Li, Q., Li, C., Yu, J., Liu, Y., Wang, G., Zhang, K., Ji, C., Yan, Q., He, L., et al.: A comprehensive survey on pretrained foundation models: A history from BERT to ChatGPT. arXiv preprint arXiv:2302.09419 (2023)