



Delft University of Technology

Document Version

Final published version

Citation (APA)

Martinez-Baselga, D., de Groot, O. M., Knödler, L., Alonso-Mora, J., Riazuelo, L., & Montano, L. (2025). Hey Robot! Personalizing Robot Navigation through Model Predictive Control with a Large Language Model. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)* (pp. 11002-11009). (Proceedings - IEEE International Conference on Robotics and Automation). IEEE. <https://doi.org/10.1109/ICRA55743.2025.11128826>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership. Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

**Green Open Access added to [TU Delft Institutional Repository](#)
as part of the Taverne amendment.**

More information about this copyright law amendment
can be found at <https://www.openaccess.nl>.

Otherwise as indicated in the copyright section:
the publisher is the copyright holder of this work and the
author uses the Dutch legislation to make this work public.

Hey Robot! Personalizing Robot Navigation through Model Predictive Control with a Large Language Model

Diego Martinez-Baselga¹, Oscar de Groot², Luzia Knoedler²,
Javier Alonso-Mora², Luis Riazuelo¹, Luis Montano¹

Abstract—Robot navigation methods allow mobile robots to operate in applications such as warehouses or hospitals. While the environment in which the robot operates imposes requirements on its navigation behavior, most existing methods do not allow the end-user to configure the robot's behavior and priorities, possibly leading to undesirable behavior (e.g., fast driving in a hospital). We propose a novel approach to adapt robot motion behavior based on natural language instructions provided by the end-user. Our zero-shot method uses an existing Visual Language Model to interpret a user text query or an image of the environment. This information is used to generate the cost function and reconfigure the parameters of a Model Predictive Controller, translating the user's instruction to the robot's motion behavior. This allows our method to safely and effectively navigate in dynamic and challenging environments. We extensively evaluate our method's individual components and demonstrate the effectiveness of our method on a ground robot in simulation and real-world experiments, and across a variety of environments and user specifications.

I. INTRODUCTION

As mobile robots increasingly operate around humans, ensuring safe and social interactions becomes essential. However, the robot's behavior requirements depend on the environment, its crowd density, and the specific robot tasks. For instance, social aspects are important requirements for elderly care, while efficiency is crucial in warehouse collaboration. A mobile robot should navigate safely in all of these scenarios, while its required tasks and behavior differ.

Recent works on socially aware navigation conclude that robots need to accommodate to human behavior and understand human intentions [1], [2], as not satisfying user preferences leads to negative experiences and frustration [3]. Behaviors can generally be configured by adapting the cost function of the motion controller. However, this requires in-depth knowledge of the system and is seldom manageable by the end-user. On the other hand, designing behaviors for all types of scenarios is unattainable. Requirements may even change over time as humans become more accustomed to

¹ The authors are with RoPeRt Group, DIIS-I3A, University of Zaragoza, 50018 Zaragoza, Spain. ({diegomartinez, riazuelo, montano}@unizar.es).

² The authors are with the Dept. of Cognitive Robotics, TU Delft, 2628 CD Delft, The Netherlands. ({o.m.degroot, l.knoedler, j.alonsomora}@tudelft.nl).

This work was supported by grant PID2022-139615OB-I00 funded by MICIU/AEI/10.13039/501100011033 and ERDF/EU, from the Aragon Government/FSE-T45_23R and from the European Union (ERC, INTER-ACT, 101041863). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

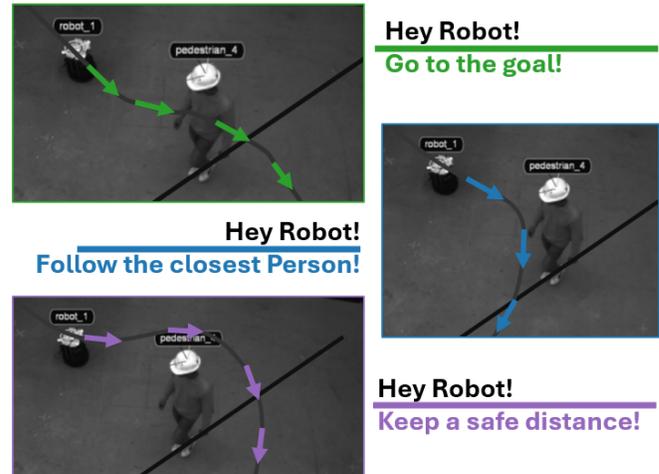


Fig. 1: A user asks the robot for a motion behavior, which is fulfilled in real-time.

the robot. Instead, the end-user should be able to modify the robot's behavior without requiring expert knowledge.

Large Language Models (LLM) and Visual Language Models (VLM) like GPT-4 [4], LLaMa [5] or Gemini [6] have emerged as powerful tools to process natural language and images and are capable of complex tasks. They pose an opportunity to interpret user requirements and requests directly from natural language. Additionally, while these methods are not trained on the specific end-user environment, they have a general understanding of real-world scenarios that can be leveraged to inform navigation behavior.

In this work, we propose an architecture composed of existing pre-trained LLM assistants and a Model Predictive Control (MPC) [7] planner to navigate challenging dynamic scenarios while accommodating user specifications, as shown in Fig. 1. Our key idea is to use the cost function of the MPC to assess and modify the capabilities of the planner in order to satisfy the user's query. Our method leverages LLMs to interpret and generate the code implementing the cost function, translating natural language instructions (e.g., "go slower", "keep more distance to people", "navigate as if you were in a hospital", "follow the path", "follow the closest human", etc.) to the desired navigation behavior. By adapting only the cost function of the MPC, we allow the user to change the robot behavior on the fly, while maintaining planner safety via constraints. As an additional feature, we equip the robot with a camera to allow it to perceive the environment independently, enabling direct adaptation based

on visual input if requested. We extensively evaluate the components of the method and the resulting navigation stack both in simulation and in hardware experiments. To this end, our contribution is three-fold:

- A novel zero-shot method to generate and tune MPC cost functions from natural language and, optionally, onboard camera images.
- A novel MPC motion planner framework for dynamic environments that may change interactively and on the fly its navigation behavior with user instructions.
- Extensive simulation and ground robot experiments that assess the performance of the proposed planner given different specified requirements.

Our implementation uses ROS, C++, Python and Acados [8] to enable real-time interactions and may be accessed in <https://github.com/dmartinezbaselga/hey-robot>.

II. RELATED WORK

Navigation in human-centered environments. Mobile robot navigation in environments shared with humans is often framed as a collision avoidance problem with dynamic obstacles. It can be approached with purely reactive methods, such as velocity obstacle-based approaches [9], [10] and the social force model [11]. Another important body of work addresses multi-agent interactions from a game-theoretic perspective [12], [13], but they present a considerable computational challenge. Reinforcement learning-based methods [14]–[16] implicitly account for interactions between agents by learning strategies through interacting with the environment. However, they lack formal guarantees, are vulnerable to out-of-distribution data, require significant training time, and produce a non-parameterized policy that cannot be easily adjusted if it fails to meet user requirements. Alternatively, MPC-based planners [7], [17], [18] explicitly account for deterministic [19] or uncertain [20], [21] future behaviors of humans when planning trajectories.

Despite major progress over the past decade, reliable robot navigation in human-centered environments remains challenging. Beyond collision avoidance, factors like comfort, naturalness, and sociability are key to human acceptance [3]. Many of the above approaches require defining desired behaviors through a cost or reward function, which demands a deep understanding of both the system and the application environment, and do not facilitate on-the-fly behavior adaptation.

LLMs in Robotics. In the last years, the application of LLMs and VLMs has revealed new opportunities for enhancing autonomous systems. VLMs hold significant potential for mobile manipulation, in particular in determining the sequence of steps needed to complete tasks [22]–[26]. These models have been employed not only to generate the necessary code for planning and executing these intermediate sub-tasks [27] but also to enable real-time, interactive communication with the robot guiding it [28]. Differently from others, [29] presents a trajectory generation algorithm

for the end-effector of a manipulator understanding language prompts with an LLM, enabling motion planning.

In mobile robots, trajectory generation from VLMs is also used to achieve zero-shot object navigation [30], [31]. In autonomous driving, for example, there are methods that incorporate them to understand the environment and output a desired car trajectory [32], directly control commands [33], [34] or select one command from a list [35]. However, the methods do not have explicit collision avoidance restrictions and have not proved to work in a closed loop in real time. As in previous works, we aim to use VLMs for high-level understanding and interactive behavior. Unlike them, we introduce the VLM inside an MPC formulation to modify the robot’s behavior in real time.

Personalized robot navigation. Recent works incorporate the gathering of human feedback to modify robot trajectory generation and use language models to process them. [36], [37] use a network to reshape trajectories given feedback and a previous trajectory, while [38] generates a costmap from a human query to be included and combined with other costs for trajectory optimization. These approaches require previous training and are for manipulators in static environments. They can not be applied to our problem, which needs fast control feedback for scenario changes.

More related to our work, some works incorporate physical interaction [39] and virtual reality [40], [41] feedback in a reinforcement learning socially-aware motion planner for mobile robots. Nevertheless, our approach is distinct from others in that it is zero-shot, uses language and visual information, and has collision avoidance constraints.

III. PROBLEM FORMULATION

We consider a robot with discrete time non-linear dynamics:

$$\mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k), \quad (1)$$

where $\mathbf{x}_k \in \mathbb{X} \subseteq \mathbb{R}^{N_x}$ is the robot state and $\mathbf{u}_k \in \mathbb{U} \subseteq \mathbb{R}^{N_u}$ the control input at time k . The state contains robot position, $\mathbf{p}_k = \{x_k, y_k\} \in \mathbb{R}^2$, and radius r_r . As the environment is dynamic, we consider $\mathcal{X} := \mathbb{R}^2 \times [0, T]$ as the robot’s workspace, with $[0, T]$ a continuous finite time domain.

The robot has to navigate in an environment with moving obstacles and avoid collisions with them and the static obstacles. We assume the moving obstacles are humans, although they could also be other robots. The humans are represented as disks with radius r_h . The position of human i at time step k is $\mathbf{o}_{i,k} \in \mathbb{R}^2$, the position of all humans is denoted \mathcal{O}_k and the area occupied by static obstacles in the environment is denoted $\mathcal{S} \subseteq \mathbb{R}^2$. The robot knows its position, the positions of the surrounding humans, their predicted trajectory, and the local costmap of the scenario. It is also equipped with a camera to capture images of its environment.

As the robot navigates, the user can provide instructions in natural language, referred to as *queries*. We denote a query j as \mathbf{q}_j , where \mathbf{q}_{j+1} is the query received after \mathbf{q}_j . Our goal is to navigate as specified by query \mathbf{q}_j and, with less emphasis,

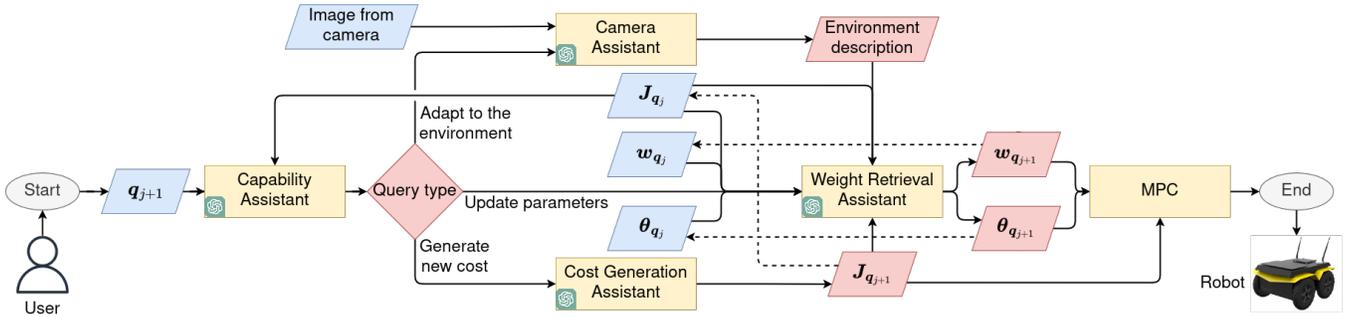


Fig. 2: Flow diagram of our proposed LLM module and its connection to the MPC. Inputs are in blue and outputs in red. Dashed lines represent updates regarding the state for the following query.

by the queries q_0, \dots, q_{j-1} as long as they do not contradict q_j . We formalize the problem as a nonlinear optimization problem over the time horizon N_k :

$$\min_{\mathbf{x} \in \mathbb{X}, \mathbf{u} \in \mathbb{U}} \sum_{k=0}^{N_k} J_{q_j}(\mathbf{x}_k, \mathbf{u}_k, \mathbf{w}_{q_j}, \boldsymbol{\theta}_{q_j}) \quad (2a)$$

$$\text{s.t.} \quad \mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k) \quad \forall k, \quad (2b)$$

$$\mathbf{x}_0 = \mathbf{x}_{init}, \quad (2c)$$

$$g(\mathbf{x}_k, \boldsymbol{\theta}_k) \leq 0 \quad \forall k, \quad (2d)$$

where in Eq. 2a the cost function that is generated to satisfy q_j is denoted J_{q_j} . The values $\mathbf{w}_{q_j} \in \mathbb{R}^{N_{w_{q_j}}}$ weight the terms in J_{q_j} (e.g., w_α weights \mathcal{J}_α), and $\boldsymbol{\theta}_{q_j}$ is a set of parameters related to the robot behavior. These include tunable parameters such as a reference velocity and environment-related parameters such as \mathcal{O}_k , \mathcal{S} , a user-specified goal, $\mathbf{p}_{goal} \in \mathbb{R}^2$, or a reference path. Equation 2b denotes the robot's dynamic constraints, Eq. 2c the initial conditions, and Eq. 2d the collision avoidance constraints. For simplicity, we model the robot dynamics with a second-order unicycle model in this paper, but the method may be applied to any robot dynamics. We consider as control input $\mathbf{u}_k = (a_k, \omega_k)$ linear acceleration and angular velocity, respectively.

IV. APPROACH

We propose Hey Robot!, a novel planning architecture (Fig. 2) that leverages LLMs to automatically generate and tune cost functions (Eq. 2) for MPC, in order to generate safe trajectories that comply with user instructions. The MPC updates control commands at every control step and is reconfigured on the fly by the LLM components when a new query, q_{j+1} , is received.

A. LLM components

The system has four different LLM components, which we call assistants, in charge of different tasks. All assistants are instructed to provide concise answers to minimize interaction time. We do ask for a brief motivation to allow the LLM to reason, improving accuracy. In this section, we use the query q_j : "Follow the path" as a running example to clarify our approach. To that end, let the cost J_{q_j} generated with respect to q_j include the following components: contour, \mathcal{J}_c , and lag, \mathcal{J}_l , terms to track a reference path in $\boldsymbol{\theta}_{q_j}$ as defined in [19];

input penalty cost terms to penalize each of the inputs $\mathcal{J}_a = a^2$, $\mathcal{J}_\omega = \omega^2$; and a velocity tracking cost, $\mathcal{J}_v = (v_k - v_{ref})^2$, to track a reference velocity, v_{ref} in $\boldsymbol{\theta}_{q_j}$: $J_{q_j} = J_{path} = w_c \mathcal{J}_c + w_l \mathcal{J}_l + w_v \mathcal{J}_v + w_a \mathcal{J}_a + w_\omega \mathcal{J}_\omega$. Our method can handle queries specifying behavior (e.g., "drive faster") and task assignments (e.g., "go to the goal").

1) *Capability Assistant*: This assistant uses LLM reasoning to decide if the current cost function J_{q_j} , written in programming code, sufficiently addresses a new user prompt q_{j+1} . We consider that J_{q_j} is sufficient for q_{j+1} if q_{j+1} can be achieved by minimizing J_{q_j} or any cost term in J_{q_j} . The query is then forwarded to the most suitable of three assistants depending on the result:

- **Generate a new cost**, if J_{q_j} is not sufficient for q_{j+1} . This implies that q_{j+1} refers to a different task like "reach the goal" or "follow a human". Then, the next assistant is the *Cost Generation Assistant*.
- **Adapt to the environment**, if q_{j+1} asks the robot to sense and adapt to its surrounding environment without any additional information (e.g. "Perceive the scenario"). The query is forwarded to *Camera Assistant*.
- **Update parameters**, if J_{q_j} is sufficient for q_{j+1} . In our example, q_{j+1} could include behavioral changes like "be faster" or "be smoother". In that case, $J_{q_{j+1}} = J_{q_j}$ and the next assistant is the *Weight Retrieval Assistant*.

2) *Cost Generation Assistant*: The goal of this assistant is to generate a new cost function $J_{q_{j+1}}$ that satisfies q_{j+1} , directly in programming language. Our approach asks the LLM to freely design $J_{q_{j+1}}$ given the query, directly writing the code of the cost function. Cost terms in $J_{q_{j+1}}$ may depend on \mathbf{x}_k , \mathbf{u}_k and parameters $\boldsymbol{\theta}_{q_{j+1}}$. Notice that through $\boldsymbol{\theta}_{q_{j+1}}$ the LLM can devise new parameters that could be updated in later queries.

To limit generation errors, the structure of the code to be generated is specified first. Additionally, a list of typical cost terms is provided for two reasons. First, we allow the LLM to name these cost terms instead of writing their code, reducing response times. Second, it may use them as code examples to generate other cost terms (few-shot prompting). The list of provided terms is \mathcal{J}_c , \mathcal{J}_l , \mathcal{J}_a , \mathcal{J}_ω , \mathcal{J}_v and a cost to reach the goal position, $\mathcal{J}_g = \|\mathbf{p}_{goal} - \mathbf{p}_k\|_2^2$. We command the LLM to always include \mathcal{J}_a , \mathcal{J}_ω and \mathcal{J}_v as part of the cost

function. If it needs to generate new cost terms different than the predefined ones, we query it to use a quadratic cost to minimize, its inverse to maximize and `if_else` CasADi [42] function to check conditions, enforcing it to use smooth and differentiable operations. The LLM generated cost function $J_{q_{j+1}}$ code is used to generate solver binaries, which are activated in the MPC. The values of $\theta_{q_{j+1}}$ and $w_{q_{j+1}}$ are tuned with the *Weight Retrieval Assistant*.

In our example, if q_{j+1} was "Reach the goal", the new cost function should be $J_{q_{j+1}} = J_{goal} = w_g \mathcal{J}_g + w_v \mathcal{J}_v + w_a \mathcal{J}_a + w_\omega \mathcal{J}_\omega$. Nevertheless, if q_{j+1} was "Follow the closest human", then a human following term should be generated from scratch, resulting in $J_{q_{j+1}} = J_{hf} = w_{hf} \|\mathbf{o}_{h,k} - \mathbf{p}_k\|_2^2 + w_v \mathcal{J}_v + w_a \mathcal{J}_a + w_\omega \mathcal{J}_\omega$, where the closest human $\mathbf{o}_{h,k} = \arg \min_{\mathbf{o}_{i,k} \in \mathcal{O}_k} \|\mathbf{o}_{i,k} - \mathbf{p}_k\|_2^2$ is retrieved using CasADi `if_else`.

3) *Camera Assistant*: This assistant uses the camera to perceive the environment and understand its implications for navigation. It allows the robot to assess, for example, if its environment is crowded, empty, narrow or open, and can recognize specific environments such as a hospital or warehouse environment. First, a photo is taken from a robot onboard camera. It is passed to the VLM, which is requested to output a list of bullet points related to ideal robot motion in the perceived scenario. We set $J_{q_{j+1}} = J_{q_j}$ to ensure that this information is solely used to inform the navigation behavior, not to change task. The description of the environment is then forwarded to the *Weight Retrieval Assistant*.

4) *Weight Retrieval Assistant*: This assistant assigns $w_{q_{j+1}}$ and tunable parameters in $\theta_{q_{j+1}}$. Its LLM prompt includes the code of $J_{q_{j+1}}$ and the values of w_{q_j} and θ_{q_j} . It also includes as query q_{j+1} or the extracted environment description when triggered by the *Camera Assistant*.

As LLMs are known to struggle with numerical tasks, it is difficult to reliably tune weights $w_{q_{j+1}}$ with an LLM. To circumvent this issue, we instead ask the LLM to rate the importance of each cost term in $J_{q_{j+1}}$ as an integer between 0 and 10, jointly denoted as $z_{q_{j+1}}$. In our example, the query "stick to the path" would result in z_c and z_l close to 10, while "be smoother" would increase the previous value of z_a and z_ω . We then convert each importance rating z_α into a weight w_α using $w_\alpha = \frac{z_\alpha}{\bar{z}_{q_{j+1}}}$, where $\bar{z}_{q_{j+1}}$ is the mean value of the elements in $z_{q_{j+1}}$. For the tunable parameters in $\theta_{q_{j+1}}$, the LLM is, in our experience, able to generate reasonable values as these parameters are absolute and have a physical meaning that the LLM understands. We therefore query the LLM directly for their values. For example, the query "be faster" would increase v_{ref} .

B. MPC formulation

The MPC controls the robot in a closed loop, solving the optimization in Eq. 2a at every control step. We use CasADi [42] and Acados [8] in the implementation. Our proposed formulation is similar to the one in [7], [17]. Having J_{q_j} , the collision constraints in Eq. 2d include human

TABLE I: Prompts tested in the *Capability, cost Generation* and *Weight retrieval assistants* experiments.

Code	Prompt
C1	Go to the goal. You are navigating through a hospital.
C2	Stick to the path.
C3	Follow the closest human.
C4	Go to the goal while keeping a safe distance from humans.
C5	Adapt to the environment.
G1	Follow the path.
G2	Reach the goal.
G3	Maximize the distance to the closest human.
G4	Minimize the distance to the closest human.
G5	Go to the goal while keeping a safe distance from humans.
G6	Follow the closest human.
W1	Be faster.
W2	Take more distance to humans.
W3	Stick to the path.
W4	Be smoother.
W5	Increase rotation capabilities.
W6	You can rotate more.

avoidance:

$$g(\mathbf{x}_k, \mathbf{o}_{j,k}) = 1 - \Delta \mathbf{p}_{j,k}^T \mathbf{R}_k^T \begin{pmatrix} \frac{1}{r^2} & 0 \\ 0 & \frac{1}{r^2} \end{pmatrix} \mathbf{R}_k \Delta \mathbf{p}_{j,k}, \quad (3)$$

where $\Delta \mathbf{p}_{j,k} = \mathbf{p}_k - \mathbf{o}_{j,k}$ and $r = r_r + r_h$. We also include static collision avoidance with \mathcal{S} , using the *Safe Flight Corridor* concept, as in [43] or [44].

MPC generally optimizes the robot trajectory from an initial guess, leading to a locally optimal trajectory that avoids the obstacles based on that guess. As the environment changes, local optimality can lead to poor and potentially dangerous navigation behavior. To compute a global optimal trajectory and avoid local minima, we use the approach from [7]. Several guidance trajectories from different homology classes [45] are sampled and optimized in parallel using Acados, resulting in locally optimal trajectories. We select the one with the minimum cost to control the robot.

V. RESULTS

We tested the individual assistants in Sec. V-A and the navigation system in simulation and on a ground robot in Sec. V-B. GPT-4o-mini [4] was used via its public API as the LLM. The MPC cost code uses Python to avoid recompilation for new costs, with C++ for static parts.

A. Assistants experiments

We evaluated the individual performance of the assistants with a series of specific queries for each, stated in Table I. The commands illustrate possible real-world user queries, e.g., C2 for reducing robot's lane departures, C5 for changing robot's operating environment, or W4 for jerky and unpredictable motion. We tried each prompt 10 times and gathered the results.

1) *Capability assistant*: The experiment uses queries in Table I as q_{j+1} and J_{path} , J_{goal} and J_{hf} as J_{q_j} . Fig. 3 represents the response rate for each q_{j+1} and J_{q_j} for different colors for the three possible responses. The expected results are that J_{path} satisfies C2, J_{goal} satisfies C1, J_{hf} satisfies C3, none of them C4 and it should adapt to the environment with C5. As the assistant has direct access to the code,

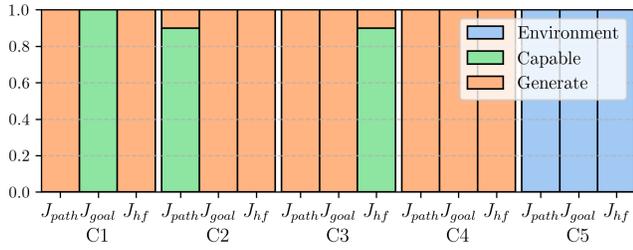


Fig. 3: Number of times (rate) each response was selected for the *Capability Assistant* experiment.

TABLE II: *Cost Generation* and *Weight Retrieval* assistants success rates.

Assistant	G1	G2	G3	G4	G5	G6
<i>Cost Generation</i>	1.0	1.0	0.9	1.0	0.9	1.0
Assistant	W1	W2	W3	W4	W5	W6
<i>Weight Retrieval</i>	1.0	1.0	1.0	0.8	1.0	1.0

written and commented using the same LLM, it successfully completes its task most of the time. The assistant only fails twice, being over-conservative, deciding to generate a new cost when it is not needed. This incurs a larger generation time, but the system still performs its task successfully.

2) *Cost Generation Assistant*: For this assistant, we verify that it reliably generates a cost function that fulfills q_{j+1} without syntax or runtime errors. We test:

- G1 and G2: Differentiating between path and goal following. In this case, it may use the names of a predefined cost terms provided in list in the prompt.
- G3 to G6: Designing by itself different cost terms that were not provided. It should design terms that maximize (G3) or minimize (G4) variables, or design high-level primitives like "keep a safe distance" (G5) or "follow a human" (G6).

Table II shows the success rates of the cost generation assistant. The assistant generates J_{path} to follow a path (G1) and J_{goal} to reach a goal (G2). It generates new cost terms that minimize (G3, G6) and maximize (G4) the distance to $o_{h,k}$ as needed, using J_{hf} and $J_{hmax} = w_{hmax} \frac{1}{(\|o_{h,k} - p_k\|_2)^2 + \epsilon} + w_v \mathcal{J}_v + w_a \mathcal{J}_a + w_\omega \mathcal{J}_\omega$, respectively. Finally, it creates a floating point parameter, the safe distance in G5, to be tuned without generating a new cost, generating $J_{sd} = w_{sd} \mathcal{J}_{sd} + w_g \mathcal{J}_g + w_v \mathcal{J}_v + w_a \mathcal{J}_a + w_\omega \mathcal{J}_\omega$, having that $d_{safe} \in \theta_{sd}$ is the safe distance; and $\mathcal{J}_{sd} = (\|o_{h,k} - p_k\|_2^2 - d_{safe})^2$ is only applied if $\|o_{h,k} - p_k\|_2^2 - d_{safe} > 0$ using the `if_else` CasADi function. The failure case in G3 was produced because it minimized instead of maximized the distance to the closest human once (we hypothesize that it is because G3 is not a common query), and in G5 because it hard coded the safe distant as a constant once, instead of including it as a parameter.

3) *Weight Retrieval Assistant*: For these experiments, we use J_{sd} to test W2 and J_{path} to test the rest of the queries. The success rates of the weight retrieval assistant are shown in Table II, computed by checking if the corresponding weights and parameters tuning correspond to the query

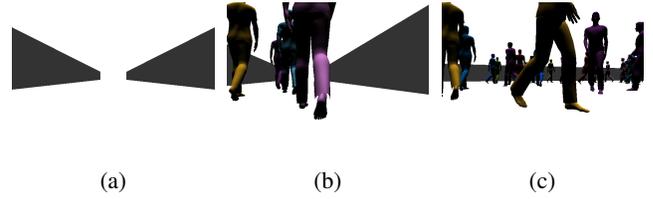


Fig. 4: Images of the simulator taken with the robot camera.

TABLE III: Final $z_{q_{j+1}}$ given the images in Fig. 4.

Image	z_c	z_l	z_v	z_a	z_ω
(a)	8	8	6	6	7
(b)	8	8	4	7	6
(c)	6	6	4	7	7

intentions. We check the final relative values of the weights to measure the results. The results show that the assistant tunes the parameters in θ_{q_j} as expected, increasing v_{ref} and d_{safe} for W1 and W2, respectively. It also increases w_v and w_{sd} with those queries, which is desirable. With W3 to W6, it proves that it understands high-level queries. It increases w_c and w_l in W3, increases w_a and w_ω in W4 and reduces w_ω in W5 and W6. The failure cases in W4 are because, as well as increasing the z_a and z_ω , it increases more z_c and z_l , so final w_a and w_ω do not increase. However, in deployment, this could be simply solved by trying the same query again or being more specific.

4) *Camera assistant*: In this experiment, we qualitatively test the *Camera Assistant* with the *Weight Retrieval Assistant* in a simulated scenario. We use J_{path} as $J_{q_{j+1}}$. We set all initial values in z_{q_j} to 5, and use the images in Fig. 4 for the queries. The results, shown in Table III, are the following:

- The scenario is described as a confined scenario with no dynamic obstacles. Thus, it sets the weights that track the reference path and reference velocity as the most important weights.
- The scenario is described as a narrow pathway with human congestion, with special requirement in careful navigation. It increases the path tracking cost due to the narrow environment but significantly reduces the velocity tracking weight, prioritizing smooth and safe navigation around humans.
- The camera detects multiple humans in close proximity and a high crowd density. As a confined room is not detected, path tracking weight is set lower, and the most important aspect is set to be a smooth navigation.

5) *Processing times*: The times taken to process the queries by the LLM in the assistant experiments presented in this section are shown in Table IV. It only takes a few seconds, which makes real-time interaction with the user possible. Regarding updates that do not require cost generation (writing code), the user may see the changes

TABLE IV: Mean and standard deviation of times spent in query processing.

Assistant	Capability	Cost Gen.	Weight Ret.	Camera
Time (s)	1.91 (1.04)	4.02 (1.63)	1.16 (0.23)	1.79 (0.38)

TABLE V: Mean and standard deviation of 10 simulations for each of six queries in the corridor-like scenario of Sec. V-B.2.

Additional Instructions	Col. rate	Dur. [s]	Path len [m]	Min Dist. [m]	v [m/s]	a [m/s ²]	ω [rad/s]
a) (Default)	0.00	14.1 (0.8)	34.99 (0.32)	0.22 (0.04)	2.49 (0.12)	0.51 (0.09)	0.18 (0.04)
b) Drive quickly.	0.00	13.7 (1.1)	34.77 (0.42)	0.19 (0.06)	2.56 (0.17)	0.59 (0.14)	0.18 (0.05)
c) Drive carefully.	0.00	24.1 (0.7)	33.30 (0.20)	0.15 (0.04)	1.38 (0.04)	0.13 (0.05)	0.05 (0.02)
d) You are navigating through a factory without humans.	0.00	13.8 (0.8)	35.03 (0.57)	0.33 (0.05)	2.55 (0.12)	0.56 (0.09)	0.20 (0.07)
e) You are navigating through a hospital.	0.00	23.6 (0.9)	33.28 (0.11)	0.25 (0.06)	1.41 (0.05)	0.10 (0.05)	0.05 (0.02)
f) Try to keep a distance of at least 1.5m from pedestrians.	0.00	19.8 (5.6)	31.09 (6.75)	1.16 (0.20)	1.74 (0.59)	0.38 (0.17)	0.16 (0.06)

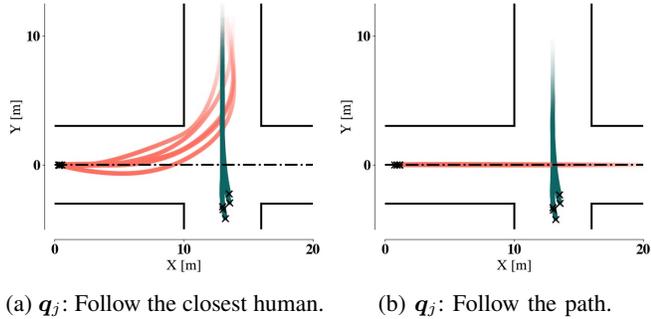


Fig. 5: Five overlapping experiments with trajectories of the robot (red) and human (blue) and a reference path (dashed).

in the robot in around 3s (*Capability Assistant* and *Weight Retrieval Assistant* times), while with queries that may change completely the behavior of the robot only around 7s (adding *Cost Generation Assistant* time).

B. Navigation Results

We tested our complete navigation stack with all its components working together with different queries, evaluating generated trajectories and metrics.

1) *Task simulations*: First, we tested our method with a simulated Clearpath Jackal robot and pedestrians moving using Social Forces [11]. We used a constant velocity assumption to predict the future trajectories of humans. First, we verify whether our method allows the user to change the robot’s task. We performed five simulations for the two queries: “follow the reference path” and “follow the closest human”. The results in Fig. 5 indicate that two *distinct* behaviors emerge, each completing the user’s task. The trajectories differ across simulations due to variations in w_j and θ_j , as the queries do not specify additional information.

2) *Navigation metrics*: We conducted ten simulations for each of six different user queries in a corridor-like scenario, with a reference path going through the center of the corridor. All queries included the task “Follow the reference path” (default), plus additional instructions. Table V shows the results. None of the experiments lead to collisions, as the MPC has collision avoidance restrictions. It may be observed that the fastest and the ones with the highest velocities are b) and d), as they are told to be fast and productive. c) and e), on the contrary, have lower velocities and are smoother, using smaller input commands. f) has the highest minimum distance to the closest pedestrian metric, as stated in the instructions, but its velocity and episode duration is not low, as other instructions were not provided. It also compromises penalizing inputs and time to reach the goal, rather than

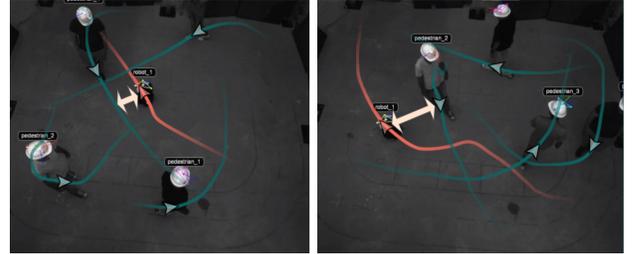


Fig. 6: Robot navigating with four humans.

the others which prefer one of them. Overall, the navigation metrics match the high-level intention of the queries.

3) *Laboratory experiments*: Moreover, we tested the ability of our method to change the robot behavior on the fly. We used a Turtlebot 2 as the ground platform, an Optitrack system with twelve cameras for localization, and a Kalman Filter with constant velocity assumption to estimate the velocity of four pedestrians present. The experiment consists of a user constantly giving different instructions to the robot while it navigates. The robot effectively addresses the user queries in successful navigation, proving the effectiveness of our solution, as seen in Fig. 6. Simulation and real-world videos of our method changing the robot behavior on the fly may be seen in the supplementary video.

VI. CONCLUSION

This work introduced a novel approach for on-the-fly personalization of mobile robot navigation using natural language. We leveraged LLMs to understand human reasoning behind user queries, and generate and tune the cost function of an MPC designed for safe navigation in dynamic environments. Our experiments were successful in both simulations and a real robot, demonstrating navigation behavior that matched the user requirements. With this work, we pave the way for new methods that reduce the gap between robots and regular users. As LLMs improve, the processing time and query understanding of our system should improve. The system is limited to the robot’s observations; queries about inaccessible elements may cause erratic behavior. Further work will include scene understanding to support object interaction, and experimentation with non-expert users to evaluate the real-world applicability and social awareness.

ACKNOWLEDGEMENTS

The use of GPT-4 public API is integrated as part of the proposed method, as explained in the manuscript.

REFERENCES

- [1] C. Mavrogiannis, F. Baldini, A. Wang, D. Zhao, P. Trautman, A. Steinfeld, and J. Oh, "Core challenges of social robot navigation: A survey," *ACM Transactions on Human-Robot Interaction*, vol. 12, no. 3, pp. 1–39, 2023.
- [2] A. Francis, C. Pérez-d'Arpino, C. Li, F. Xia, A. Alahi, R. Alami, A. Bera, A. Biswas, J. Biswas, R. Chandra, *et al.*, "Principles and guidelines for evaluating social robot navigation algorithms," *arXiv preprint arXiv:2306.16740*, 2023.
- [3] T. Kruse, A. K. Pandey, R. Alami, and A. Kirsch, "Human-aware robot navigation: A survey," *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1726–1743, 2013.
- [4] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, *et al.*, "Gpt-4 technical report," *arXiv preprint arXiv:2303.08774*, 2023.
- [5] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, *et al.*, "Llama: Open and efficient foundation language models," *arXiv preprint arXiv:2302.13971*, 2023.
- [6] M. Reid, N. Savinov, D. Teplyashin, D. Lepikhin, T. Lillicrap, J.-b. Alayrac, R. Soricut, A. Lazaridou, O. Firat, J. Schrittwieser, *et al.*, "Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context," *arXiv preprint arXiv:2403.05530*, 2024.
- [7] O. de Groot, L. Ferranti, D. Gavrila, and J. Alonso-Mora, "Topology-driven parallel trajectory optimization in dynamic environments," *arXiv preprint arXiv:2401.06021*, 2024.
- [8] R. Verschuere, G. Frison, D. Kouzoupis, J. Frey, N. van Duijkeren, A. Zanelli, B. Novoselnik, T. Albin, R. Quirynen, and M. Diehl, "acados – a modular open-source framework for fast embedded optimal control," *Mathematical Programming Computation*, 2021.
- [9] J. Alonso-Mora, P. Beardsley, and R. Siegwart, "Cooperative collision avoidance for nonholonomic robots," *IEEE Transactions on Robotics*, vol. 34, no. 2, pp. 404–420, 2018.
- [10] D. Martínez-Baselga, E. Sebastián, E. Montijano, L. Riazuelo, C. Sagiúés, and L. Montano, "Avocado: Adaptive optimal collision avoidance driven by opinion," *arXiv preprint arXiv:2407.00507*, 2024.
- [11] D. Helbing and P. Molnar, "Social force model for pedestrian dynamics," *Physical review E*, vol. 51, no. 5, p. 4282, 1995.
- [12] L. Peters, A. Bajcsy, C.-Y. Chiu, D. Fridovich-Keil, F. Laine, L. Ferranti, and J. Alonso-Mora, "Contingency games for multi-agent interaction," *IEEE Robotics and Automation Letters*, 2024.
- [13] N. Mehr, M. Wang, M. Bhatt, and M. Schwager, "Maximum-entropy multi-agent dynamic games: Forward and inverse solutions," *IEEE transactions on robotics*, vol. 39, no. 3, pp. 1801–1815, 2023.
- [14] C. Chen, Y. Liu, S. Kreiss, and A. Alahi, "Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning," in *2019 international conference on robotics and automation (ICRA)*. IEEE, 2019, pp. 6015–6022.
- [15] M. Everett, Y. F. Chen, and J. P. How, "Collision avoidance in pedestrian-rich environments with deep reinforcement learning," *Ieee Access*, vol. 9, pp. 10357–10377, 2021.
- [16] D. Martínez-Baselga, L. Riazuelo, and L. Montano, "Improving robot navigation in crowded environments using intrinsic rewards," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 9428–9434.
- [17] O. de Groot, L. Ferranti, D. Gavrila, and J. Alonso-Mora, "Globally guided trajectory planning in dynamic environments," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 10118–10124.
- [18] D. Martínez-Baselga, O. de Groot, L. Knoedler, L. Riazuelo, J. Alonso-Mora, and L. Montano, "Shine: Social homology identification for navigation in crowded environments," *arXiv preprint arXiv:2404.16705*, 2024.
- [19] B. Brito, B. Floor, L. Ferranti, and J. Alonso-Mora, "Model predictive contouring control for collision avoidance in unstructured dynamic environments," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 4459–4466, 2019.
- [20] H. Zhu and J. Alonso-Mora, "Chance-constrained collision avoidance for mavs in dynamic environments," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 776–783, 2019.
- [21] A. Wang, X. Huang, A. Jasour, and B. Williams, "Fast risk assessment for autonomous vehicles using learned models of agent futures," *arXiv preprint arXiv:2005.13458*, 2020.
- [22] W. Huang, P. Abbeel, D. Pathak, and I. Mordatch, "Language models as zero-shot planners: Extracting actionable knowledge for embodied agents," in *International conference on machine learning*. PMLR, 2022, pp. 9118–9147.
- [23] A. Brohan, Y. Chebotar, C. Finn, K. Hausman, A. Herzog, D. Ho, J. Ibarz, A. Irpan, E. Jang, R. Julian, *et al.*, "Do as i can, not as i say: Grounding language in robotic affordances," in *Conference on robot learning*. PMLR, 2023, pp. 287–318.
- [24] H. Liu, Y. Zhu, K. Kato, A. Tsukahara, I. Kondo, T. Aoyama, and Y. Hasegawa, "Enhancing the llm-based robot manipulation through human-robot collaboration," *arXiv preprint arXiv:2406.14097*, 2024.
- [25] Y. Jin, D. Li, A. Yong, J. Shi, P. Hao, F. Sun, J. Zhang, and B. Fang, "Robotgpt: Robot manipulation learning from chatgpt," *IEEE Robotics and Automation Letters*, 2024.
- [26] C. Winge, A. Imdieke, B. Aldeeb, D. Kang, and K. Desingh, "Talk through it: End user directed manipulation learning," *arXiv preprint arXiv:2402.12509*, 2024.
- [27] J. Liang, W. Huang, F. Xia, P. Xu, K. Hausman, B. Ichter, P. Florence, and A. Zeng, "Code as policies: Language model programs for embodied control," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 9493–9500.
- [28] C. Lynch, A. Wahid, J. Tompson, T. Ding, J. Betker, R. Baruch, T. Armstrong, and P. Florence, "Interactive language: Talking to robots in real time," *IEEE Robotics and Automation Letters*, 2023.
- [29] T. Kwon, N. Di Palo, and E. Johns, "Language models as zero-shot trajectory generators," *IEEE Robotics and Automation Letters*, 2024.
- [30] A. Majumdar, G. Aggarwal, B. Devnani, J. Hoffman, and D. Batra, "Zson: Zero-shot object-goal navigation using multimodal goal embeddings," *Advances in Neural Information Processing Systems*, vol. 35, pp. 32340–32352, 2022.
- [31] Y. Dai, R. Peng, S. Li, and J. Chai, "Think, act, and ask: Open-world interactive personalized robot navigation," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 3296–3303.
- [32] C. Sima, K. Renz, K. Chitta, L. Chen, H. Zhang, C. Xie, P. Luo, A. Geiger, and H. Li, "Drivelm: Driving with graph visual question answering," *arXiv preprint arXiv:2312.14150*, 2023.
- [33] H. Shao, Y. Hu, L. Wang, G. Song, S. L. Waslander, Y. Liu, and H. Li, "Lmdrive: Closed-loop end-to-end driving with large language models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 15120–15130.
- [34] L. Chen, O. Sinavski, J. Hünemann, A. Karnsund, A. J. Willmott, D. Birch, D. Maund, and J. Shotton, "Driving with llms: Fusing object-level vector modality for explainable autonomous driving," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 14093–14100.
- [35] H. Sha, Y. Mu, Y. Jiang, L. Chen, C. Xu, P. Luo, S. E. Li, M. Tomizuka, W. Zhan, and M. Ding, "Languagempc: Large language models as decision makers for autonomous driving," *arXiv preprint arXiv:2310.03026*, 2023.
- [36] A. Bucker, L. Figueredo, S. Haddadin, A. Kapoor, S. Ma, and R. Bonatti, "Reshaping robot trajectories using natural language commands: A study of multi-modal data alignment using transformers," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 978–984.
- [37] A. Bucker, L. Figueredo, S. Haddadin, A. Kapoor, S. Ma, S. Vemprala, and R. Bonatti, "Latte: Language trajectory transformer," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 7287–7294.
- [38] P. Sharma, B. Sundaralingam, V. Blukis, C. Paxton, T. Hermans, A. Torralba, J. Andreas, and D. Fox, "Correcting Robot Plans with Natural Language Feedback," in *Proceedings of Robotics: Science and Systems*, New York City, NY, USA, June 2022.
- [39] M. Kollmitz, T. Koller, J. Boedecker, and W. Burgard, "Learning human-aware robot navigation from physical interaction via inverse reinforcement learning," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 11025–11031.
- [40] J. de Heuvel, N. Corral, L. Bruckschen, and M. Bennewitz, "Learning personalized human-aware robot navigation using virtual reality demonstrations from a user study," in *2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 2022, pp. 898–905.
- [41] J. de Heuvel, N. Corral, B. Kreis, J. Conradi, A. Driemel, and M. Bennewitz, "Learning depth vision-based personalized robot navigation

- from dynamic demonstrations in virtual reality,” in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 6757–6764.
- [42] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, “CasADi – A software framework for nonlinear optimization and optimal control,” *Mathematical Programming Computation*, 2018.
- [43] S. Liu, M. Watterson, K. Mohta, K. Sun, S. Bhattacharya, C. J. Taylor, and V. Kumar, “Planning dynamically feasible trajectories for quadrotors using safe flight corridors in 3-d complex environments,” *IEEE Robotics and Automation Letters*, vol. 2, no. 3, pp. 1688–1695, 2017.
- [44] D. Benders, J. Köhler, T. Niesten, R. Babuška, J. Alonso-Mora, and L. Ferranti, “Embedded hierarchical mpc for autonomous navigation,” *arXiv preprint arXiv:2406.11506*, 2024.
- [45] S. Bhattacharya, M. Likhachev, and V. Kumar, “Topological constraints in search-based robot path planning,” *Autonomous Robots*, vol. 33, pp. 273–290, 2012.