Gesture2Path: Imitation Learning for Gesture-aware Navigation

Catie Cuan^{1,2}, Edward Lee³, Emre Fisher², Anthony Francis³, Leila Takayama³, Tingnan Zhang³, Alexander Toshev^{4,†}, and Sören Pirk^{5,†}

Abstract—As robots increasingly enter human-centered environments, they must not only be able to navigate safely around humans, but also adhere to complex social norms. Humans often rely on non-verbal communication through gestures and facial expressions when navigating around other people, especially in densely occupied spaces. Consequently, robots also need to be able to interpret gestures as part of solving social navigation tasks. To this end, we present Gesture2Path, a novel social navigation approach that combines image-based imitation learning with model-predictive control. Gestures are interpreted based on a neural network that operates on streams of images, while we use a state-of-the-art model predictive control algorithm to solve point-to-point navigation tasks. We deploy our method on real robots and showcase the effectiveness of our approach for the four gestures-navigation scenarios: left/right, follow me, and make a circle. Our experiments indicate that our method is able to successfully interpret complex human gestures and to use them as a signal to generate socially compliant trajectories for navigation tasks. We validated our method based on in-situ ratings of participants interacting with the robots.

I. INTRODUCTION

Situated agents should not only navigate safely around people, but should also abide by social norms and respond to the full gamut of human behavior - including nonverbal communication such as body language, gestures, and facial expressions. A robot solving a navigation task must be able to interpret human behavior and to carefully adjust its actions to be socially compliant. We refer to this form of navigation task as Nonverbal Social Navigation. For defining and generating social behavior, robotics research has recently expanded efforts toward understanding the importance of respecting personal space [1] and social dynamics [21], [37], socially-acceptable behavior for approaching humans [18], navigation among groups of people [22], the validation of socially acceptable policies [31], and curating large datasets [20]. The breadth of these research directions is a testament of the complexity of generating socially compliant agent behavior.

To solve navigation tasks, many existing approaches rely on point cloud data obtained from LiDAR scanners that provide real-time information about the environment, including dynamic objects such as humans. The captured point

³Robotics at Google, 1600 Amphitheatre Parkway, Mountain View, CA 94304, USA

⁵Adobe Research, 345 Park Avenue, San Jose, CA 95110-2704, USA



Fig. 1: *Gesture2Path* is a novel social navigation policy that combines image-based imitation learning with model-predictive control to enable gesture-aware navigation. In this example we show our *right* gesture policy: A robot at its start location begins navigating toward its goal (a). On its way, it encounters a person that indicates to the robot with a *right* gesture to pass on their right (b). The robot interprets this gesture and drives around the person in the intended manner (c) to then continue its path towards its goal (d). The robot in use is from Everyday Robots.

clouds are projected to 2-D occupancy maps. With these maps existing policy algorithms – such as model predictive control (MPC) – can efficiently solve complex navigation tasks with remarkable success. However, while point cloud data is a powerful sensor modality for generating navigation trajectories in the environment, it rarely provides the accuracy needed to interpret intricate human behavior. Conversely, while RGB sensors do not provide the depth information needed to solve navigation tasks, they produce high-resolution images that allow the capture of nuanced gestures and facial expressions.

In this paper, we propose a novel gesture-aware navigation policy based on imitation learning and MPC. We train a sequential neural network that enables us to generate waypoints of navigation trajectories from a sequence of consecutive images. Our goal is to train this network so that it predicts waypoints that are socially compliant and adherent to the gestures of humans interacting with the robot. Once the network is trained, the predicted waypoints are sent to an MPC algorithm to control the robot. The sequential neural network serves as a high-level planner, while the MPC algorithm provides low-level control. This setup provides the benefits of both approaches: the sequential neural network allows us to obtain nuanced information of human gestures from sequences of images, while the MPC algorithm allows us to safely navigate the robot based on laser-scanned point clouds.

¹Department of Mechanical Engineering, Stanford University, Stanford, 94305, United States

²Everyday Robots, 100 Mayfield Ave, Palo Alto, CA 94043, USA

⁴Apple ML Research, Apple Park 1, Cupertino, CA 95014, USA

[†] Work done while at Robotics at Google.



Fig. 2: Overview of *Gesture2Path*: A sequential neural network trained via imitation learning acts as a high-level planner to predict waypoint deltas from image sequences. The predicted waypoint deltas are used to compute a world space waypoint that serves as intermediate target position. MPC uses these intermediate target positions to compute linear and angular velocities to drive the robot.

To explore the effectiveness of our gesture-aware social navigation policy, we define the four gesture scenarios: Left/Right, Follow Me, and Make a Circle. For each scenario we define the gesture (e.g. pointing with the hand to the left) and the robot behavior (e.g. drive left) so each gesture maps to a specific robot behavior – an example is shown in Fig. 1. We then collected a dataset of expert examples for each scenario by driving the robot with a human operator. A collected trajectory is defined by a sequence of images and a sequence of robot waypoints. We then train the sequential neural network on subsequences of images to predict subsequences of future waypoints. The MPC algorithm then computes the linear and angular velocities to drive from robots' current position to the predicted position of the sequential model.

We validate our gesture-aware social navigation policy based on in-situ experience ratings of humans interacting with the robots. To obtain these ratings we follow a protocol for social navigation policies [31]. Each of our scenarios is defined by a gesture as well as the start and end positions of the trajectory that the robot is supposed to drive along, mirroring the training setup. Additionally, we define a questionnaire based on a five-level Likert scale for each scenario. Once a scenario is completed a participant is asked to provide their ratings for a specific scenario. We use this setup to compare our *Gesture2Path* policy with a MPC policy as the baseline.

In summary, our contributions are: (1) we introduce *Ges*ture2Path, a novel gesture-aware social navigation policy that combines a sequential neural network with MPC; (2) we show that our policy is able to interpret gestures to generate sociallycompliant behavior in gesture-navigation scenarios; (3) we define a canonical set of social navigation scenarios with gestures that can be replicated and tested; (4) we employ a novel validation protocol for social navigation scenarios to compare our gesture-aware policy with a standard MPC policy.

II. PRIOR WORK

Due to its importance to robotics, research on social navigation has recently garnered considerable attention, spanning a breadth of methods which we cannot comprehensively discuss. For a general overview, interested readers are referred to recent survey papers by Gao and Huang [14], Charalampous et al. [6], Mavrogiannis et al. [27], Kruse et al. [25], Rios-Martinez [32], and Mirsky et al. [28]. The following discussion focuses on approaches to social navigation and human robot interaction with methods closely related to our work. A variety of researchers have explored gestures for robotic control. For teleoperation, researchers used hand tracking with a Microsoft Kinect device [39] [10], a Leap motion [38], or a camera to direct a robot's end effector. For robot navigation, [33], [34] extracted gestures from a Kinect sensor to generate navigation commands for a teleoperated robot in a non-social setting. [19] used hidden Markov models to detect six gestures with a Kinect sensor and mapped them to hand-crafted robot navigation behaviors. [5] used a wearable gesture detector to control a mobile robot. Finally, [3], [2] and [26] used imitation learning to replicate human gestures on a robot platform.

Researchers have also examined gestures for social communication. [17] modeled musical improvisation as a series of gestures so a robot marimba player could improvise with a human. In [29], researchers designed a gesture framework for humans to collaborate with a co-worker robot on an assembly task; [11] approached a different collaborative task between humans and robots using a gesture framework. Other gesture-based human-robot communication systems have been developed for settings where verbal communication may be challenging, such as emergency settings [8] or underwater [7].

III. METHOD

Our main goal is to develop a gesture-aware navigation policy which can guide robots around humans in a socially compliant manner while responding to provided gestures. We define our Gesture2Path policy by combining a sequential neural network as a high-level planner with MPC for low level control (Fig. 2). The high-level planner identifies humans and the gestures they perform from images, while the lowlevel controller safely navigates the robot from one waypoint to another. The advantage of this setup is that the high-level planner and the low-level controller complement each other. MPC performs well for navigating from one waypoint to another, even around dynamic obstacles, based on point cloud data obtained from the environment. However, MPC fails to detect humans and their gestures as it does not use RGB data. Conversely, the high-level planning neural network is able to detect humans and their gestures from sequences of images, but fails to reliably predict waypoints for the navigation task.

A. Gesture-aware Social Navigation

To control a robot in response to a human and their gestures, we employ a sequential neural network [35] that, given a history of observations, outputs a sequence of future positions. We learn this model from demonstrations of successful navigation trajectories of human-robot interactions and show that the network is able to learn to interpret human gestures.

In more detail, we denote by s = (I, O, p) the robot state consisting of an RGB image I from a headmounted camera, an occupancy grid O from a LiDAR sensor, and a robot position p in the world coordinate frame. A demonstration trajectory of length n can be described as a sequence of states: (s_1, \ldots, s_n) . An occupancy map is defined as the 2D projection of a point cloud obtained from a LiDAR scanner.

Given a history of states, we aim to compute controls that would result in the future robot positions as encoded by the demonstrations. In particular, at step n we utilize a history of k states from the demonstration to predict l future positions:

$$(p_{n+1}, \dots, p_{n+l}) = \text{Gesture2Path}(s_n, \dots, s_{n-k}, g), \quad (1)$$

where g is the goal position of the robot.

In theory, we could use a neural network to directly learn robot control commands. Several approaches do learn direct mappings from pixels to motor controls [23], [24], [30]. However, learning an end-to-end robot control policy remains challenging, as methods tend to require large training datasets and easily overfit to the underlying robot dynamics.

Therefore, we employ MPC as a low level controller that, given an intermediate goal position, drives the robot to this position while avoiding collisions. The MPC algorithm takes care of converting target positions to low level torque controls.

To navigate a robot a MPC controller requires the intermediate next position encoded in the robot frame. Thus, we need to convert the demonstrated robot positions from world frame to a sequence of egocentric positions, each in the frame of the previous position. This can be done by calculating the deltas between subsequent positions: $\delta_i = p_i - p_{i-1}$ for $i = 1, \ldots, l$, which leads to the final model formulation:

$$(\delta_{n+1},\ldots,\delta_{n+l}) = \text{Gesture2Path}(s_n,\ldots,s_{n-k},g).$$
(2)

1) Sequential Model Architecture: We train the sequential model by simultaneously embedding subsequences of RGB images I, the occupancy maps O, and the target deltas $au = g - p_n$ into a joined latent space. Specifically, our neural network architecture uses 3D convolutional layers to obtain two 64-dimensional embeddings from RGB images and occupancy maps of the input subsequence with a resolution of 256x256 pixels. To combine the target deltas (τ) with the image embeddings we also project them into a 64-dimensional latent space. We then add the embeddings and pass them to three dense layers to obtain a sequence of l position deltas (δ). For the last dense layer we use a linear activation for regression. Fig. 3 shows this network architecture. Compared to other approaches for action recognition (e.g. [15]), the goal for defining a lightweight multi-modal architecture for detecting gestures is to ensure small inference times (< 30ms).

While our policy is agnostic to the choice of the sequential model, we found a common 3D-convolutional neural network similar to [4] provided the best performance for our experiments. We train the network via imitation learning using expert sequences of human-robot interactions captured by manually driving the robot with a human operator (see Section IV-B).



Fig. 3: We use a 3D CovNet neural network architecture that joinly embeds sequences of RGB images I, occupancy maps O, and target deltas τ into a fused embedding space. The network is trained so as to predict sequences of position deltas that we use to generate an intermediate goal position for an MPC algorithm.



Fig. 4: Our sequential neural network predicts a subsequence of position deltas (δ) from a history of past states *s*. We use the predicted position deltas to compute an intermediate target position (p_t).

Fig. 4 illustrates the parameters of our learning setup. We train our sequential neural network on subsequences of the collected expert trajectories (full sequence). The training objective is to predict a sequence of position deltas (with l steps into the future) based on a history of k previous states. The goal is to use the predicted position deltas to compute an intermediate target position (p_t) from the current position (p_n) . To train the model we sample all possible subsequences from our dataset of expert trajectories.

2) Action Parameterization: We use MPC for low-level robot control. Specifically, we formulate the goal following as an optimization problem and use an iterative linear quadratic regulator (iLQR) solver, Trajax [13], that minimizes a handengineered cost function, similar to [9], [16], [36]. The optimization problem was solved at each control step and the predicted action at the first step (i.e. velocity command) is sent to the robot. The inputs to MPC include the occupancy grid and relative goal position; its cost function combines time-weighted goal penalties, margin-offset collision penalties, and a weighted control penalty; and its action space is linear and angular velocities. This MPC policy provides smooth navigation and fast reaction times for differential drive robots and is superior in our tests to a re-implementation of the wellperforming reinforcement learning policy in [12].

We use MPC in two ways: first we establish a baseline, where we use the algorithm to navigate from a start to a goal position without additional visual inputs. Second, we use MPC for our gesture-aware policy by tracking waypoints that are generated by the sequential neural network. In both cases, the action space are linear and angular velocities.

For gesture-aware navigation, we use the sequential neural network described in Section III-A1 to predict position deltas



Fig. 5: Once the intermediate target position p_t has been computed from the predicted position deltas, we use MPC to compute an optimal trajectory from the current position p_n to the target position. Note that the trajectory represented by the predicted position deltas and the computed MPC trajectory may diverge – we do not want to over-constrain the MPC algorithm to find an optimal trajectory.



Fig. 6: We test *Gesture2Path* with three policies to respond to four gestures: circle (a), follow (b), right (c), and left (d). For each scenario we only define the gesture (e.g. pointing with the hand to the left), the corresponding expected robot behavior (e.g. drive left), and the start and end positions of the robot trajectory. The robot is tasked to navigate from its start position to a goal area (blue circle). A person is standing in the center of the room and interacts with the robot by performing different gestures to initiate the robots' response.

from a sequence of previous states. From the position deltas (δ) we compute an intermediate target position $p_t = p_n + \sum_{i=n+1}^{i=n+1} \delta_i$ for MPC based on the current position p_n . The MPC algorithm then computes an optimal path from the current position p_n to the intermediate target position p_t .

At every time step we collect the robot state s_n . To generate gesture-aware trajectories we use the neural network to make predictions every k time steps (Fig. 2, Planning Loop), while the MPC algorithm runs at a frequency of 10 Hz (Fig. 2, Control Loop). With k = 10, this means that the neural network is used to predict a new intermediate target position every 10 time steps, while the MPC algorithm continuously generates control commands to drive toward the intermediate target position. Note that the number of previous states (defined by k) used to make predictions with the neural network need not be the same as the number of predicted position deltas (defined by l). The neural network may be triggered to predict a new intermediate target position before the robot reaches the previously computed one. This ensures that the gesture-aware policy is susceptible to chances in the observed gestures.

Moreover, as illustrated in Fig. 5, the trajectory represented by the predicted position deltas of the sequential neural network and the MPC-computed trajectory may diverge. While we want the robot to closely follow the predicted position deltas to implement gesture-aware behavior, we do not want to over-constrain MPC to find an optimal trajectory from the current position to the intermediate target position. The planning loop based on the sequential neural network provides general direction while considering observed gestures, and the control loop based on MPC ensures safe navigation.

IV. EXPERIMENT SETUP

The robots we use are 7 DOF mobile manipulators with a single arm and a rectangular base. Each robot is equipped with a head-mounted camera and a LiDAR sensor. Robots also have access to an existing static map of their environment to assist with localization and point-to-point navigation.

A. Gesture Scenarios

To test the capabilities of *Gesture2Path*, we use the protocol proposed in [31] to define the four gesture scenarios (illustrated in Fig. 6): Circle, Follow, and Left/Right

- **Circle**: In the make a circle scenario, the demonstrator would stand in the center of the room and raise their right hand straight up to the ceiling. The robot operator then drives the robot to make a counterclockwise circle before passing the human on their left (Fig. 6 a).
- Follow: In the follow me scenario, the demonstrator would lift both hands out to the side while keeping their feet together, making a "T" shape. After holding this shape for 2 seconds, they would turn over their left shoulder and walk along a wide arc prior to passing through the end position region. Once the gesture has been observed the robot operator would then drive the robot so as to follow the demonstrator (Fig. 6 b).
- Left/Right: In the left/right scenario, the demonstrator would stand in the center of the room and put their arm out to the left or right while leaning in that direction. The robot operator would then drive the robot to pass along the side of the human's outstretched arm (Fig. 6 c, d).

We selected these gestures based on the following criteria for the human and the robot: (1) the robot solves a navigation tast beginning at one side of the room and ending at the other; (2) the robot would move continuously from beginning to end without pausing or stopping; (3) the gesture instructions would be simple and easy to teach the experiment participants; (4) the gestures would correspond to clear changes in robot behavior; (5) the gestures require the robot to react to a single participant. While the gestures used for the Circle and Follow scenarios are distinct, the Left and Right gestures represent more subtle variations for the sequential model to disentangle.

B. Data Collection and Training

To train our *Gesture2Path* policy we collected a dataset of 277 trajectories (40 for follow, 30 for circle, 207 left and right). Each trajectory represents a 20-40 second sequence that includes a human-robot interaction. The human performs one of the four defined gestures, while the robot solves a point-to-point navigation task from a start to a goal position. All training data was collected in the real world in a large, open atrium space over several different days and times (Fig. 7).

To collect our dataset we used the following procedure: an expert robot operator would connect to the real robot using a Logitech F710 gamepad. The robot operator would stand



Fig. 7: For data collection the robot is controlled by a joystick handled by an experienced teleoperator. The human demonstrator stands in the center of the room and performs the gesture as well as the next sequence of human motions (if in the follow condition). The teleoperator moves the robot from the beginning position to the ending position with the correct response.

behind the robot and follow it while driving it forward. The gesture demonstrator would begin by standing in the middle of the room. Once the teleoperated robot started to move, the gesture demonstrator would perform the gesture for 2-3 seconds. Depending on the gesture scenario, the demonstrator would either stay in place or move appropriately.

We split the collected trajectories into 90% training and 10% validation data. We then trained our sequential neural network on subsequences (k = 10, l = 10, for most of our experiments) of RGB images and occupancy maps. We used the Adam optimizer with a standard l2 loss and selected checkpoints of the network for our policies based on the lowest l2 error measured on the validation dataset. The training of our network usually converges after 120 epochs, which usually takes 6 hours of training on a single V100 GPU.

C. Experiment Design

To validate the effectiveness of our gesture-aware policies we conducted a user-study to obtain in-situ experience ratings. In the following, we discuss the setup of our experiment for our *Gesture2Path* polices as well as for the MPC baseline.

1) User Study Design: We designed a within-subjects experiment where participants performed the four gestures while the robot was running either *Gesture2Path* or a baseline of the unmodified MPC policy. Note this baseline MPC policy is itself a highly performant navigation policy, exceeding in our internal tests a reimplementation of [12]. The study occurred in the same location where the training data was collected.

Participants arrived at the location and filled out an introductory survey asking about their dominant hand and number of months they had been working near these robots. The participants then performed each of the four gestures a total of 70 different times in the center of the room. Each the trials consisted of the robot crossing from one side of the room to the other; the human participants would perform the gesture in the center of the room (and during the follow scenario, walking to the edge of the room). Once the trial was completed the participant would immediately answer four questions about the robot's performance. The questionnaire was defined based on a 5-level Likert scale. Each of the four questions were rated on a scale of 1-5, 1 being *Strongly Disagree* and 5 being *Strongly Agree*. Participants had the option to answer 0 if the question did not apply. The questions were:

- 1) The robot maintained a safe distance at all times. (Safe)
- 2) It was clear what the robot wanted to do. (Clarity)
- 3) The robot responded correctly to my gesture. (Correct)
- 4) The robot paid attention to what I was doing. (Trust)



Fig. 8: The means and standard errors for all four questions for policy and gesture condition. The MPC conditions are in grey and *Gesture2Path* policy in blue. Participants rated the question on a scale of 1 Strongly Disagree to 5 Strongly Agree. Zero was reserved for Not Applicable. The policy rated overall more highly for each of the four questions when compared to MPC. A paired t-test was performed on each MPC/policy bar for each gesture (i.e. comparing circle MPC to circle policy). All of these tests were statistically significant with a p-value < .01. We can conclude that the policy performs overall better than MPC for each of these gesture scenarios.

We ran the gesture trials 10 times for each of our *Gesture2path* policies (Follow, Circle, Left/Right). For the MPC baseline we ran 10 trials for the Follow and Circle scenarios and 5 trials for the Left and Right scenarios.

We randomized policy ordering for each participant. For example, one participant may have completed MPC Left/Right, *Gesture2Path* Left/Right, MPC Follow, Policy Follow, MPC Circle, Follow Circle. After the gesture trials, the participants filled out a closing survey asking about their age, gender, cultural group, height, highest degree of education received, primary field of work or study, as well as their experience with STEM, robotics in general, and specifically the used robot. The experiment took approx. 90 minutes to complete.

2) Participants: 12 individuals volunteered to participate in the user study, including 2 women and 10 men. 8 participants identified as Asian, 2 as White, 1 as Hispanic/Latino, and 1 Middle Eastern. 3 participants were ages 18-24, 8 were ages 25-34, and 1 declined to state. All participants listed their *right hand* as the dominant hand and all participants had previously used the robot in their work.

V. RESULTS

All participant ratings for each gesture were grouped together by *Gesture2Path* and MPC (i.e. questions 1-4 for Circle gesture, *Gesture2Path*, were all grouped). These grouped responses were averaged, as shown in Fig. 9. A paired t-test was performed for each gesture vs policy group (i.e. Questions 1-4 for Circle *Gesture2Path* vs Questions 1-4 for Circle MPC) and all differences are statistically significant with a p-value < .01. Therefore, we conclude that the *Gesture2Path* policy overall performs better than MPC for all four gesture scenarios.

To determine which questions had the strongest effect on *Gesture2Path's* improvement over MPC, we ran a repeated



Fig. 9: All averages and standard errors for each question for each gesture, *Gesture2Path* is in blue, MPC is in grey. Question 3 (Correctness) shows the largest difference between *Gesture2Path* and MPC. A repeated measures ANOVA on Question 3 showed a statistically significant difference.

measures ANOVA on each question. The repeated measures ANOVA comparing the effect of *Gesture2Path* vs MPC on Question 3 (Correctness) showed a statistically significant difference: F(within groups df) = 115.563, p-value = .059.

Separate repeated measures ANOVAs were performed to compare the effect of *Gesture2Path* vs MPC on Question 1 (Safety), Question 2 (Clarity) and Questions 4 (Trustworthiness). Fig. 8 shows that *Gesture2Path* (blue bars) is higher in all cases for Question 2 (Clarity) and Question 4 (Trustworthiness), while MPC (grey bars) is higher for Question 1 (Safety). However, for Questions 1, 2, and 4, the ANOVA did not show statistical significance. Therefore, the critical improvement that elevated *Gesture2Path* over MPC was the policy's correct response to the human gesture (Question 3).

The *Gesture2Path* policies took longer than MPC in all cases, as summarized in Table I. For Circle and Follow, this was likely due to the gesture response trajectory requiring a longer distance for the robot to travel. For the Left/Right Policy, the additional time was likely due to the wider distance needed by the gesture, as well as how quickly the human operator drove the robot during data collection. A paired t-test was performed on each Gesture2Path and MPC pairing for each gesture and the results were statistically significant.

In Fig. 11 we show 2-dimensional plots of the robot navigating from the start to end positions. The positive participant ratings (4 or 5 out of 5) for Question 3 (Correctness) are shown in solid magenta. The neutral rating (3 out of 5) is show in dotted black. The negative ratings (1 or 2 out of 5) are shown in dotted red. These plots demonstrate that the participants have a clear understanding of the robot's correct response as the higher ratings correspond with the robot's correct behavior.

Finally, in Fig. 12, we show examples for the four scenarios. For each policy the image sequences show (from left to right) the robot at its start location, the human performing the



Fig. 10: Example of the MPC policy: As no image sequences are used, the policy is not able to generate trajectories in response to gestures. While the performed gesture is supposed to make the robot drive a circle, it only passes the human participant on their right side.

TABLE I: p-values for gesture, MPC vs. policy, and relevant question using a paired t-test

Gesture	Gesture2Path, MPC	Avg duration (sec)	p-value
Circle	Gesture2Path, MPC	56, 35	< .001
Follow	Gesture2Path, MPC	45, 38	< .01
Left	Gesture2Path, MPC	45, 38	< .01
Right	Gesture2Path, MPC	46, 33	< .01

The Gesture2Path policy took significantly longer than MPC in all cases, likely due to the increased distance requirements of the correct response path or the speed of the demonstrations.

gesture, the gesture response of the robot, and the navigation to the goal location. As shown, our *Gesture2Path* policy is able to generate socially compliant trajectories in response to the performed gestures. In contrast, as the MPC policy is not using image sequences as input to detect gestures, it is not able to generate trajectories based on the performed gestures – an example is shown in Fig. 10.

A. Discussion and Limitations

Our MPC baseline can reliably navigate around obstacles with success comparable to published state of the art [12]. In this work we show that using imitation learning in tandem with MPC simultaneously allows us to maintain this performance level while enabling gesture-aware navigation. The aim of the study was to show that robots can respond correctly to human gestures and we observe this in the results.

The trained *Gesture2Path* policies perform better than baseline MPC policy for all gestures in the participant ratings (Fig. 9). Question 3 (Correctness) was statistically significant, showing that *Gesture2Path* is able to correctly identify and respond to human gestures. The difference in ratings between baseline MPC and *Gesture2Path* was larger for the Circle and Follow policy than for the Right and Left policy, meaning that either participants were more certain of the robot's correct response or that *Gesture2Path*'s behavior was more distinctive than baseline MPC. Baseline MPC performed better than *Gesture2Path* for all four gestures on Question 1 (Safety). The higher safety rating follows from baseline MPC's tendency to give humans a large amount of space to avoid collisions.

Recall we trained three policies: 1) Circle only, 2) Follow only, and 3) Left and Right. The third policy must generate different paths for two gestures. As baseline MPC tends to go right when it encounters obstacles, it appears correct to participants half the time (high clarity rating, Question 2).

We limited the number of questions to four due to practicality reasons; running the study became quite lengthy. The study ran for a minimum of 90 minutes, so participants may have been fatigued towards the end. One participant



Fig. 11: 2D plots of the robot trajectories and corresponding participant ratings (Question 3 only) for all four gestures with Gesture2Path (Ours) and MPC. The robot starts in the red circle and is tasked to navigate to the blue dot. The human is the black circle in the center. For the Follow plots, there is an additional black alternating dot-dash line to indicate the rough path that the human walks after they make the gesture in the center of the room. Favorable participant ratings (4 or 5 on a 5 point scale) are solid lines in magenta. Neutral participant ratings (3 on a 5 point scale) are the dashed line in blue. Negative participant ratings (1 or 2 on a 5 point scale) are the dotted lines in red. The MPC trajectories in the first row are all similar regardless of the gesture. The Gesture2Path policy robot trajectories in b) Circle and d) Follow are distinctive, correct, and consistent, with participant ratings capturing this effect. The Gesture2Path policy does not perform as well for Right and Left sides, also reflected in the participant ratings.

elected not to complete the study halfway through due to the time commitment. To mitigate this fatigue effect in the data, conditions were randomized for each participant.

Furthermore, not all participants perform the gesture the same way. Despite watching the same video in order to learn the gesture, each human has different timings, range of motion, and qualitative emphasis. Therefore, it is promising that *Gesture2Path* can generalize to these performance variations.

VI. CONCLUSION AND FUTURE WORK

We have introduced Gesture2Path, a novel social navigation policy, to enable gesture-aware point-to-point navigation. Our approach combines a novel sequential neural network which can predict position deltas from images with MPC for generating linear and angular velocities to control a robot. As the sequential neural network is trained on images of the observed gestures it can generate positions as waypoints which implement socially-acceptable and gesture-aware navigation behavior. On the other hand, MPC allows us to safely navigate from the current position of the robot to an intermediate goal position. Based on four scenarios we have shown that our Gesture2Path policy is able to respond correctly, clearly, and legibly to a range of human gestures, while it also outperforms the baseline MPC algorithm devoid of visual inputs.

In future work, we plan to train our policy so as to handle multiple gestures simultaneously. This step is important for deploying Gesture2Path in a working robot environment, where the aim is to maximize the number of gestures with the smallest number of concurrently running policies. We will also explore multiple humans gesturing simultaneously. As another direction for future work we are interested in exploring a closer connection of the introduced sequential neural network and MPC. While keeping both algorithms separate has intriguing advantages from an engineering point of view, it seems likely

that a closer integration of understanding gestures and lowlevel control would lead to better performance. Finally, we want to use our Gesture2Path policy for learning even more complex and diverse sets of gestures to enable even more advanced robot navigation behavior.

REFERENCES

- [1] P. Althaus, H. Ishiguro, T. Kanda, T. Miyashita, and H.I. Christensen. Navigation for human-robot interaction tasks. In IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA '04. 2004, volume 2, pages 1894-1900 Vol.2, 2004.
- [2] J.P. et. al Bandera-Rubio. Vision-based gesture recognition in a robot learning by imitation framework. 2010.
- [3] S. Calinon and A. Billard. Learning of gestures by imitation in a humanoid robot. Technical report, Cambridge University Press, 2007.
- [4] J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. pages 4724–4733, 07 2017. D.J. Chaithanya, D. Gowda, D. Balakrishna, D. M. Gowda, and MS
- [5] Divyashree. A survey on hand gesture controlled rover. 2022.
- [6] K. Charalampous, I. Kostavelis, and A. Gasteratos. Recent trends in social aware robot navigation: A survey. Robotics and Autonomous Systems, 93:85-104, 2017.
- [7] D. Chiarella, M. Bibuli, G. Bruzzone, M. Caccia, A. Ranieri, E. Zereik, L. Marconi, and P. Cutugno. A novel gesture-based language for underwater human-robot interaction. Journal of Marine Science and Engineering, 6(3):91, 2018.
- [8] F. De Cillis, G. Oliva, F. Pascucci, R. Setola, and M. Tesei. On field gesture-based human-robot interface for emergency responders. In 2013 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR), pages 1-6. IEEE, 2013.
- [9] J. Di Carlo, P. M. Wensing, B. Katz, G. Bledt, and S. Kim. Dynamic locomotion in the mit cheetah 3 through convex model-predictive control. In *IROS*, pages 1–9, 2018. [10] G. Du, P. Zhang, J. Mai, and Z. Li. Markerless kinect-based hand
- tracking for robot teleoperation. International Journal of Advanced Robotic Systems, 9(2):36, 2012.
- [11] T. Ende, S. Haddadin, S. Parusel, T. Wüsthoff, M. Hassenzahl, and A. Albu-Schäffer. A human-centered approach to robot gesture based communication within collaborative working processes. In 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 3367-3374. IEEE, 2011.
- [12] A. Francis, A. Faust, L. Chiang, J. Hsu, J. C. Kew, M. Fiser, and E. Lee. Long-range indoor navigation with prm-rl. IEEE Transactions on Robotics, 36(4):1115-1134, 2020.
- [13] R. Frostig, V. Sindhwani, S.T. Singh, and S.N. Tu. trajax: differentiable optimal control on accelerators. 2021.



Fig. 12: The *Gesture2Path* policy employed for the four gesture-aware navigation scenarios: circle (a), follow (b), left (c), and right (d). For all scenarios the robots starts on the right side of the room. The robot in the photo is from Everyday Robots.

- [14] Y. Gao and C.-M. Huang. Evaluation of socially-aware robot navigation. *Frontiers in Robotics and AI*, 8, 2022.
- [15] R. Girdhar, M. Singh, N. Ravi, L. van der Maaten, A. Joulin, and I. Misra. Omnivore: A single model for many visual modalities. 2022.
- [16] M. Hertneck, J. Köhler, S. Trimpe, and F. Allgöwer. Learning an approximate model predictive controller with guarantees. *IEEE Control Systems Letters*, 2(3):543–548, 2018.
- [17] G. Hoffman and G. Weinberg. Gesture-based human-robot jazz improvisation. In *ICRA*, pages 582–587. IEEE, 2010.
- [18] C.-M. Huang, T. Iio, S. Satake, and T. Kanda. Modeling and controlling friendliness for an interactive museum robot. 07 2014.
- [19] L. Iocchi and M.T. Bonanni. Person-tracking and gesture-driven interaction with a mobile robot using the kinect sensor. 2011.
 [20] H. Karnan, A. Nair, X. Xiao, G. Warnell, S. Pirk, A. Toshev, J. Hart,
- [20] H. Karnan, A. Nair, X. Xiao, G. Warnell, S. Pirk, A. Toshev, J. Hart, J. Biswas, and P. Stone. Socially compliant navigation dataset (scand): A large-scale dataset of demonstrations for social navigation. *IEEE Robotics and Automation Letters*, pages 1–8, 01 2022.
- [21] Y. Kato, T. Kanda, and H. Ishiguro. May i help you? design of humanlike polite approaching behavior. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, HRI '15, page 35–42, New York, NY, USA, 2015. Association for Computing Machinery.
- [22] K. D. Katyal, K. Popek, G. D. Hager, I.-J. Wang, and C.-M. Huang. Prediction-based uncertainty estimation for adaptive crowd navigation. In *Artificial Intelligence in HCI (AI-HCI)*, page 353–368, Berlin, Heidelberg, 2020. Springer-Verlag.
- [23] H. Kim, M. Jordan, S. Sastry, and A. Ng. Autonomous helicopter flight via reinforcement learning. In S. Thrun, L. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems*, volume 16. MIT Press, 2003.
- [24] N. Kohl and P. Stone. Policy gradient reinforcement learning for fast quadrupedal locomotion. In *IEEE International Conference on Robotics* and Automation, 2004. Proceedings. ICRA '04. 2004, volume 3, pages 2619–2624 Vol.3, 2004.
- [25] T. Kruse, A. K. Pandey, R. Alami, and A. Kirsch. Human-aware robot navigation: A survey. *Robotics and Autonomous Systems*, 61(12):1726– 1743, 2013.
- [26] Y. Kuniyoshi. Learning from examples: Imitation learning and emerging cognition. *Humanoid robotics and neuroscience: science, engineering* and society, pages 234–249, 2015.

- [27] C. Mavrogiannis, F. Baldini, A. Wang, D. Zhao, P. Trautman, A. Steinfeld, and J. Oh. Core challenges of social robot navigation: A survey, 2021.
- [28] R. Mirsky, X. Xiao, J. Hart, and P. Stone. Prevention and resolution of conflicts in social navigation–a survey. arXiv preprint arXiv:2106.12113, 2021.
- [29] P. Neto, M. Simão, N. Mendes, and M. Safeea. Gesture-based humanrobot interaction for human assistance in manufacturing. *The International Journal of Advanced Manufacturing Technology*, 101(1):119–135, 2019.
- [30] J. Peters and S. Schaal. Reinforcement learning of motor skills with policy gradients. *Neural Networks*, 21(4):682–697, 2008. Robotics and Neuroscience.
- [31] S. Pirk, E. Lee, X. Xiao, L. Takayama, A. Francis, and A. Toshev. A protocol for validating social navigation policies. *ICRA, Workshop: Social Robot Navigation: Advances and Evaluation*, 2022.
- [32] J. Rios-Martinez, A. Spalanzani, and C. Laugier. From proxemics theory to socially-aware navigation: A survey. *International Journal of Social Robotics*, 7(2):137–153, Apr 2015.
- [33] S. Saha, R. Lahiri, A. Konar, A. Lekova, and A.K. Nagar. A novel approach to tsk model based gesture driven robot movement. In 2017 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE), pages 1–6. IEEE, 2017.
- [34] S. Saha, R. Lahiri, A. Konar, and A.K. Nagar. Gesture driven remote robot manoeuvre for unstructured environment. In 2018 IEEE Symposium Series on Computational Intelligence (SSCI), pages 1793–1800. IEEE, 2018.
- [35] I. Sutskever, O. Vinyals, and Q. V. Le. Sequence to sequence learning with neural networks. *NIPS*, 27, 2014.
- [36] Y. Tassa, N. Mansard, and E. Todorov. Control-limited differential dynamic programming. In *ICRA*, pages 1168–1175, 2014.
 [37] X.-T. Truong and T.-D. Ngo. "to approach humans?": A unified
- [37] X.-T. Truong and T.-D. Ngo. "to approach humans?": A unified framework for approaching pose prediction and socially aware robot navigation. *IEEE Transactions on Cognitive and Developmental Systems*, 10(3):557–572, 2018.
- [38] W. Zhang, H. Cheng, L. Zhao, L. Hao, M. Tao, and C. Xiang. A gesture-based teleoperation system for compliant robot motion. *Applied Sciences*, 9(24):5290, 2019.
- [39] Z. Zhang. Microsoft kinect sensor and its effect. *IEEE multimedia*, 19(2):4–10, 2012.