



**HAL**  
open science

# DeepMoTion: Learning to Navigate Like Humans

Mahmoud Hamandi, Mike d'Arcy, Pooyan Fazli

► **To cite this version:**

Mahmoud Hamandi, Mike d'Arcy, Pooyan Fazli. DeepMoTion: Learning to Navigate Like Humans. 28th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN 2019), Oct 2019, New Delhi, India. pp.1-7, 10.1109/RO-MAN46459.2019.8956408 . hal-04863066

**HAL Id: hal-04863066**

**<https://laas.hal.science/hal-04863066v1>**

Submitted on 3 Jan 2025

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# DeepMoTion: Learning to Navigate Like Humans

Mahmoud Hamandi<sup>1</sup>, Mike D’Arcy<sup>2</sup>, and Pooyan Fazli<sup>3</sup>

**Abstract**—We present a novel human-aware navigation approach, where the robot learns to mimic humans to navigate safely in crowds. The presented model, referred to as DeepMoTion, is trained with pedestrian surveillance data to predict human velocity in the environment. The robot processes LiDAR scans via the trained network to navigate to the target location. We conduct extensive experiments to assess the components of our network and prove their necessity to imitate humans. Our experiments show that DeepMoTion outperforms all the benchmarks in terms of human imitation, achieving a 24% reduction in time series-based path deviation over the next best approach. In addition, while many other approaches often failed to reach the target, our method reached the target in 100% of the test cases while complying with social norms and ensuring human safety.

## I. INTRODUCTION

Robots are gradually moving from factories and labs to streets, homes, offices, and healthcare facilities. These robots are currently assigned tasks that require interaction with humans, such as guiding passengers through busy airport terminals [20] or roaming around university buildings and interacting with nearby humans [15], [23].

As robots are increasingly becoming part of our everyday lives, it is essential for them to be aware of the surrounding humans while performing their tasks. Navigation is a basic skill for autonomous robots, but many traditional algorithms, such as A\* and D\*, do not consider the fact that the obstacles in the environment may be humans. While maneuvers made by these algorithms may produce short paths and avoid direct collisions, they do not consider social norms, such as walking on the right side and passing on the left. This can cause inconvenience for humans. We define human-aware navigation as the ability of the robot to navigate while complying with social norms and ensuring human safety.

While many existing systems allow robots to navigate safely within crowds [6], [19], they still rely heavily on manually crafted models of human motion. Such models may capture the aspects of human motion as understood by their designers, while they may likely miss subtle trends that characterize their human aspect. In addition, manually crafted models do not have a way to automatically adapt to different cultures, so it may require significant manual effort to be used in a different environment.

We present DeepMoTion (Deep Model for Target-driven Imitation), a deep imitation learning algorithm that eliminates the need for an explicit model of human motion and instead learns the human navigation patterns directly by observing pedestrians. By imitating the motion patterns learned from real humans, the algorithm naturally follows social norms without needing such rules to be manually specified. Moreover, the network learns to decide on the direction and speed associated with raw LiDAR data without any preprocessing. The network is trained to learn the possible motion patterns it might face in human crowds on its own.

The goal of this work is to learn directly from pedestrian data without the need for a predefined human model. With the absence of a true model, learning the reward governing human motion is not feasible with current Inverse Reinforcement Learning algorithms such as the one presented in [24]. Our method tackles the imitation problem as a classification one, where the network learns a specific command for each observation without simulating the learned policy. This approach reduces the amount of time required for each architecture test and allows us to explore multiple network configurations.

The contributions of the paper are as follows:

- 1) We present a deep imitation learning algorithm to generate navigational commands and plan a path to the target in the environment, similar to humans, for a mobile ground robot. The proposed method outperforms all the benchmarks on time series-based path deviation and reaches the target in 100% of the test cases while complying with social norms and ensuring human safety.
- 2) Further, we present a novel loss function to train the network. The loss function allows us to accommodate for human motion stochasticity while at the same time enabling the robot to navigate safely.

We conduct extensive experiments to assess the components of our deep neural network and prove their necessity to imitate humans.

## II. BACKGROUND AND STATE OF THE ART

Previous work on human-aware navigation suggested to apply handcrafted models to control a robot about humans, define human-centric cost maps, or even follow humans through crowds.

Helbing and Molar [9] presented the Social Force Model (SFM), where they modeled the assumed social forces governing the human motion. Ferrer *et al.* [6] used the social force model to navigate in a way similar to humans. In their

<sup>1</sup>Mahmoud Hamandi is with LAAS-CNRS, Université de Toulouse, CNRS, Toulouse, France [mhamandi@laas.fr](mailto:mhamandi@laas.fr)

<sup>2</sup>Mike D’Arcy is with the Department of Computer Science, Northwestern University, Evanston, IL, USA [m.m.darcy@u.northwestern.edu](mailto:m.m.darcy@u.northwestern.edu)

<sup>3</sup>Pooyan Fazli is with the Department of Computer Science, San Francisco State University, San Francisco, CA, USA [pooyan@sfsu.edu](mailto:pooyan@sfsu.edu)

work, the robot navigates to the target while abiding by the social forces, that is, the robot is attracted by its target and repelled by pedestrians and obstacles. Furthermore, they extended the social force model to allow the robot to escort a human while providing a scheme to learn the parameters of the model.

Inverse reinforcement learning (IRL) has also been investigated to learn human-like navigation policies from data. Henry *et al.* [10] adapted the MaxEnt IRL algorithm to partially-observed environments for socially-aware navigation. Vasquez *et al.* [21] tested a variety of features, such as crowd density and the social forces to learn a cost map that replicates the reward maximized in human navigation. Kim and Pineau [12] developed a navigation system based on maximum a posteriori Bayesian IRL. Using IRL for human-like navigation is a two-step process: first, training a reward function from human data using IRL, and then using a separate algorithm to find actions for the robot that are optimal under the learned reward function. On the other hand, in DeepMoTion, the network is directly trained to produce actions that lead to human-like navigation.

Bera *et al.* [2] predicted human motion after observing a set of related psychological cues, such as aggressiveness, tension, and level of activity. The robot then deduces the acceptable path from the predicted human locations and social distances inferred from the same psychological features.

Another approach to human-aware navigation was presented by Mehta *et al.* [14], where the robot follows a human through crowds when it cannot navigate on its own. In their approach, the robot decides to navigate freely when the scene is clear or compromise its optimal shortest path by following a human to its goal. When neither possibility is viable, the robot stops and waits for a clearance.

Sisbot *et al.* [19] suggested a set of human-centric costs that allow the robot to navigate safely around humans. The method applied a cost-based navigation algorithm with a Gaussian coercing a safety distance about each human. In addition, the robot attempted to stay in the visual range of the existing pedestrians and to increase its own visibility near hidden areas, such as when rotating around a corner.

Lu and Smart [13] proposed another method where the robot navigates following a human-aware cost map. Their approach forces the robot to navigate on the right side of a hallway, allowing opposing humans to navigate on its left. In addition, the robot communicated its awareness of the nearby pedestrians by tilting its head toward their eyes.

While these methods provide a model for human motion, multiple deep learning architectures were also presented in the literature to learn any navigation algorithm. Pfeiffer *et al.* [17] proposed an end-to-end network that allowed the robot to navigate based on LiDAR scans and target position. Similarly, Groshev *et al.* [8] presented a network that learns reactive policies that imitate a planning algorithm when provided with current and goal observations. Both papers presented novel ideas, however, they learn reactive policies ignoring previous robot states while trying to imitate long-term planning algorithms.

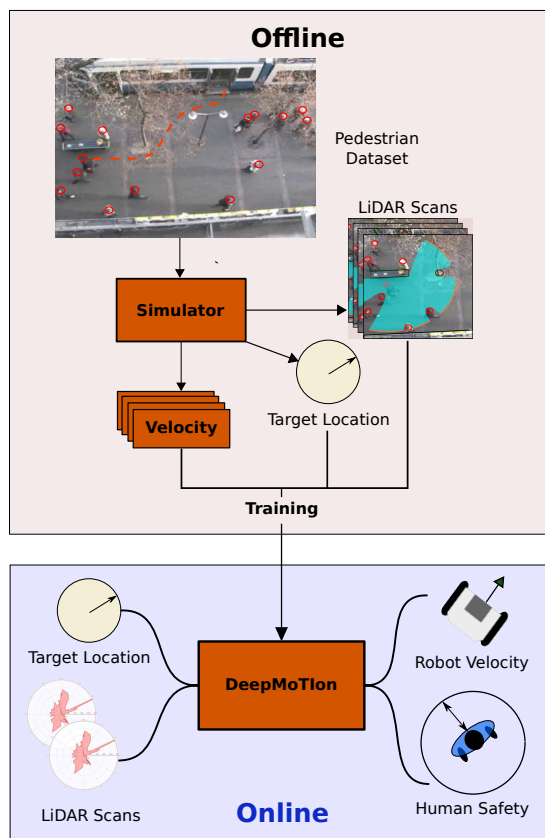


Fig. 1: Algorithm overview: In the offline phase, we train a deep neural network, called DeepMoTion, based on a pedestrian dataset. In the online phase, the network is given the target location and the last two consecutive LiDAR scans of the environment, and the network produces speed and direction outputs that are safe and adhere to social norms.

Chen *et al.* [3] propose a socially-aware navigation approach using deep reinforcement learning. However, the socially-aware behavior of the method was achieved using a system of handcrafted rewards when training the model. This contrasts with our method, which aims to learn a socially-aware navigation policy directly from human trajectory data without manually specifying social norms, such as passing on the left.

Crowd simulation methods such as [4] and [11] aim to produce realistic simulations of human movement. However, while the methods use local policies that can be conditioned on a desired destination, they use handcrafted models of human motion with just a few tunable parameters.

We contrast our problem on human motion imitation with the literature on human motion prediction [1], [5], where the primary focus is predicting a human’s future position based on their current location, surroundings, and a history of previous states. The predictions in these methods are not conditioned on a known target, making them unsuitable for the navigation task. With human motion imitation, we assume the robot is given a target location, and its objective is to determine how a human would navigate to it.

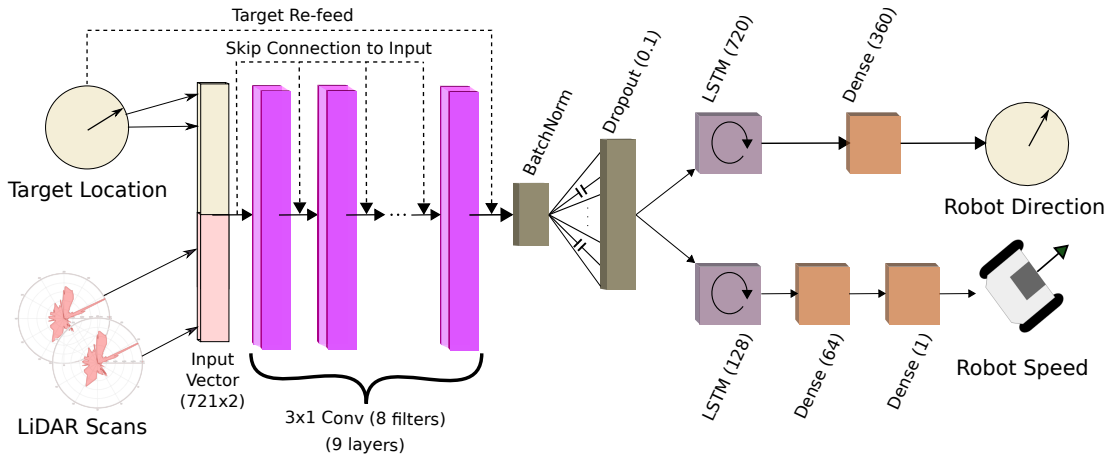


Fig. 2: DeepMoTion network architecture. DeepMoTion processes the target location and the last two consecutive LiDAR scans via a series of convolutional layers with skip connections followed by batch normalization and dropout. The result is then processed by two separate branches to produce the direction output and the speed output. Each branch consists of LSTM and dense layers.

In conclusion, methods in the literature show the advantage of imitating humans, although they do not do so from actual human traces. In what follows, we intend to train directly from pedestrian data to alleviate the need for any human modeling through an end-to-end network.

### III. PROBLEM DEFINITION

Our assumption is that the best way to teach a robot to navigate is to let it learn directly from observing humans' navigation. In this approach, we replace humans one at a time in a pedestrian dataset with our robot equipped with a limited range  $360^\circ$  LiDAR sensor and let it observe the environment at each time step. Then the robot should learn to mimic the human's navigation for the given observation.

Figure 1 shows the different parts of our human motion imitation method. In the offline phase, we train a deep neural network, called DeepMoTion, based on a pedestrian dataset. In the online phase, the network is given the target direction, target distance, and the last two consecutive LiDAR scans of the environment, and in return the network provides the robot with navigational commands to reach the target while moving similar to humans and ensuring human safety.

Our model is trained with the ETH pedestrian dataset [16] presenting videos of humans navigating in a real-world environment. The dataset contains environment maps and a set  $\chi$  of humans, and for each human  $h$  the trajectory  $\zeta_h$  that they took through the environment. Each  $\zeta_h$  is a sequence of locations  $l_{h,t}$ , representing the position of human  $h$  at time  $t$ . We use a simulator to estimate the target location  $\tau_{h,t}$ , LiDAR scan  $\mathbf{z}_{h,t}$ , and velocity  $\mathbf{v}_{h,t}$  at every time step  $t$  for each human replaced by the robot in the dataset, which we then use to train the network to imitate the human trajectories. The simulator uses a manually constructed static obstacle map for each environment along with the annotated human trajectories to approximate the LiDAR scans.

After training, the robot processes its target location and the last two consecutive LiDAR scans via DeepMoTion

to calculate navigational commands that allow it to reach the target safely while moving similarly to the humans in the dataset. This end-to-end learning happens only through observing humans' navigation. The ETH pedestrian dataset is challenging for any autonomous robot due to the dense crowds and sudden changes of pedestrians' directions.

### IV. DEEPMOTION

DeepMoTion is a deep neural network  $f(s_{h,t})$  defined as:

$$\mathbf{v}_{h,t} = (d_{h,t}, v_{h,t}) = f(s_{h,t}),$$

$$s_{h,t} = \begin{bmatrix} \mathbf{z}_{h,t-1} & \tau_{h,t} \\ \mathbf{z}_{h,t} & \tau_{h,t} \end{bmatrix}, \quad (1)$$

where  $s_{h,t}$  is the input state matrix, and  $(d_{h,t}, v_{h,t})$  is the output action set describing the direction and magnitude of the velocity  $\mathbf{v}_{h,t}$ . The output direction  $d_{h,t}$  is represented as a 360-dimensional vector. This is converted to a scalar heading in degrees by taking the argmax. The input target location  $\tau_{h,t}$  is represented in a similar way, being formed by concatenating a 360-dimensional one-hot target direction vector with a scalar target distance to produce a 361-dimensional target location vector. Note that the direction vectors for the input and output can have any number of dimensions, as this simply controls the angular resolution of the direction. We use 360-dimensional vectors because they provide reasonably precise angular resolution and have a convenient translation to degrees.

As shown in Equation (1), our network receives the current and previous LiDAR scans in addition to the target location. We found that providing the two LiDAR scans greatly improved the performance of the network compared to only giving the current LiDAR scan. With only one LiDAR scan, the network had difficulty distinguishing between moving and static obstacles.

Our deep neural network architecture is shown in Figure 2. The input to the network is the target location ( $361 \times 1$ ) and




Fig. 3: Gaussian distribution (blue) about the human-chosen direction (red) with a standard deviation  $\sigma$ .

the last two consecutive LiDAR observations ( $360 \times 2$ ). We concatenate two copies of the target location ( $361 \times 2$ ) with the  $360 \times 2$  LiDAR matrix to create the network’s input vector ( $721 \times 2$ ). The network has 9 convolutional layers, each with 8 ( $3 \times 1$ ) filters and a stride size of 1. The input to each filter was padded to conserve its size, and each filter was followed by a *tanh* activation function. In this architecture, the skip connections from the input were inspired by classical planning algorithms [8], such as value iteration and greedy search. However, after the convolutional layers, we re-feed only the raw target location to the network due to its direct correlation to the velocity direction, while the LiDAR scans add minimal value in their raw state. We found through experimentation that only shared convolutional layers were required for the network to correctly deduce the direction and speed from the input state, while adding specialized convolutional layers for each of the two outputs, similar to [8], reduced its performance.

In addition, for a planning algorithm each state and the corresponding action are tightly related to the previous observations. The LSTM layer was added to the network to keep some memory of all the previous steps, because these layers have been shown to improve the prediction of future states based on their memory of the past [7]. We later provide a thorough experimental comparison to show the LSTM’s necessity. Batch normalization was necessary to assure the boundedness of the input to the LSTM layers. The final dense layers process the LSTM output to provide the direction and speed to be used by the robot.

#### A. Loss Function

Our loss function is designed to train the network to output the direction and speed as seen in the human dataset by minimizing the squared error of the speed and the cross entropy error of the output direction.

However, human imitation presents a challenge due to its stochasticity. In fact, two humans might behave differently even with the same observations depending on their personality and other hidden factors. This suggests that the correct direction might be one of many directions in a range about the ground-truth. As such, it is desirable to penalize

the network less for cases where it is close to the ground truth than cases where it is completely wrong. To this end, we model the output direction as a Gaussian distribution about the human-chosen direction with a standard deviation  $\sigma$  shown in Figure 3. The standard deviation  $\sigma$  over the human chosen direction is challenging to derive from the pedestrian dataset due to the fact that we only observe a single sample for any given navigational situation. We do not know the range of possible directions that different humans in the same situation might choose because the chosen direction is conditioned on the target position, current location, and observed obstacles, which are different at each timestep for each human. Therefore, we cannot simply compute the standard deviation over all direction choices made by humans in the dataset, and we instead tune  $\sigma$  as a model hyperparameter.

As such, our complete loss function for a batch of  $N$  training examples can be expressed as follows:

$$\overbrace{\frac{1}{N} \sum_{i=1}^N (v_i - \hat{v}_i)^2}^{\text{Speed Loss}} + \overbrace{\frac{1}{N} \sum_{i=1}^N H(d_i(\sigma), \hat{d}_i)}^{\text{Direction Loss}}, \quad (2)$$

where  $v_i$  is the actual human speed,  $\hat{v}_i$  is the predicted speed,  $H(\cdot, \cdot)$  is the cross entropy loss function,  $d_i(\sigma)$  is the Gaussian distribution about the human-chosen direction with standard deviation  $\sigma$ , and  $\hat{d}_i$  is the predicted direction distribution.

## V. EXPERIMENTS AND RESULTS

To evaluate the performance of DeepMoTion when imitating humans, we conducted experiments on the ETH BIWI walking pedestrians dataset [16]. The dataset provides annotated trajectories of 650 humans recorded over 25 minutes of time on two maps. We randomly assigned 2/3 of the data to the training set and 1/3 to the test set. To avoid overfitting and allow the network to generalize to unseen maps, the training data was augmented by replicating each path while rotating the map at random angles.

The dense crowds and sudden changes of pedestrian direction in this dataset make it sufficiently challenging for our experiments. We conducted two types of experiments on this dataset. First, we assessed the different components of the network, and then we compared it with other benchmark methods in terms of human imitation, safety, and target reachability.

It should be noted that our algorithm runs in real-time despite the depth of the network due to the low dimensionality of its input state vector. GPU was only used for training. During testing, the algorithm runs in real-time on a single core of a 3.3 GHz CPU (0.084 seconds per forward pass), which can be set up easily on a mobile ground robot.

#### A. Benchmarks

To assess the performance of DeepMoTion, we ran several baseline algorithms in our experiments for comparison.

These include two deep learning algorithms as well as a human-aware navigation method from the literature.

- 1) **Generalized Reactive Planner (GRP):** GRP is a deep neural network architecture composed of multiple convolutional layers, each with a skip connection from the input, followed by fully connected layers to the output [8]. In addition, the input of GRP is a concatenation of the observation and target. GRP is trained to learn reactive policies that allow the robot to imitate a planning algorithm.
- 2) **End-to-end Motion Planning (EMP):** EMP is another deep neural architecture that relies on a relatively small number of deep convolutional layers and two residual shortcut connections [17]. Moreover, the target in this network is provided after the convolutional layers. EMP is trained to learn a navigation algorithm.
- 3) **Social Force Model (SFM):** SFM calculates a set of imaginary ‘social’ forces that govern the human motion in a crowd [6]. These forces can be grouped as repulsive to obstacles and other humans as well as attractive to the target.

The optimizer for all deep architectures in our experiments (GRP, EMP, and DeepMoTion) was Adadelta, with a learning rate of 1 and an L2 regularization weight of 0.001.

### B. Metrics

To assess our network and the benchmark algorithms, we compared their performance when trying to navigate from a start position to the final target. The start and final positions are chosen from the dataset, where the simulator replaces one of the humans with a robot and compares the resemblance of their paths, as well as the safety of the robot and other humans in the environment. Formally, we compare the performance of our network and the benchmark algorithms based on the following metrics:

- 1) **Squared Path Difference (SPD):** The trajectories of the robot and the corresponding human are modeled as discrete-time trajectories  $T_{r,0..n}$  and  $T_{h,0..m}$  respectively. The squared path difference can then be expressed as

$$\sum_{i=1}^{\max(n,m)} \|T_{r,i} - T_{h,i}\|^2, \quad (3)$$

where the last location in the shorter path is compared with the remaining steps of the other. This metric indirectly penalizes the difference in length between the robot’s and the human’s paths.

- 2) **Dynamic Time Warping (DTW):** *DTW* is a metric described in [18] to measure the similarity between two temporal sequences, which may vary in speed. The metric finds the optimal time warp to match the segments of the two paths and measures the similarity between them following that warp. While *SPD* reflects an algorithm’s ability to replicate both direction and speed, *DTW* compares the two paths irrespective of their speed. For instance, similarities in walking could

be detected using *DTW*, even if the imitating robot was walking faster or slower than the human, or if there were accelerations and decelerations over the course of the navigation.

- 3) **Proximity:** *Proximity* is the closest distance the robot comes to a human on its path. In the case of any collision along the path, it is assigned a value of 0. We report the average *proximity* over all the test cases.
- 4) **Number of Collisions:** The number of times the robot collides with a human while navigating in the environment.
- 5) **Target:** The percentage of trials where the robot reaches the goal within the 400-step threshold.

It should be noted that unlike most human-aware navigation papers [6], [22], we are reporting the average number of collisions as a comparison metric. However, when implemented in real-world settings, a low-level obstacle avoidance controller is to be added to the algorithm to assure complete human safety and accommodate for any failures similar to [12].

### C. DeepMoTion Variants

To study the necessity of the LSTM layer, we tested two variants of the network, where we refer to DeepMoTion<sub>LSTM</sub> as the network with the architecture explained before and DeepMoTion<sub>conv</sub> as the network without any LSTM layers. More convolutional layers were added to DeepMoTion<sub>conv</sub> to accommodate for the depth difference. Both variants of DeepMoTion were trained with the full loss function shown in Equation (2) and with a fixed  $\sigma = 5$ .

Table I shows that DeepMoTion<sub>LSTM</sub> imitates humans better than DeepMoTion<sub>conv</sub> and exhibits a better performance in all metrics. DeepMoTion<sub>LSTM</sub> outperforms DeepMoTion<sub>conv</sub> with regard to path difference (*SPD* and *DTW*) and safety (*proximity* and the number of collisions) and reaches the target in 100% of the trials. The results prove the necessity of the LSTM layer and the fact that having two consecutive LiDAR scans fed into the network does not work on its own.

Figure 4 also shows example robots navigating with the two variants of DeepMoTion’s architecture. Figure 4 (left) shows an example where DeepMoTion<sub>conv</sub> is accumulating error throughout its path and finally misses the target. This behavior was observed throughout the trials on many occasions, which explains the difference in performance of the two networks. The figure also shows that DeepMoTion<sub>LSTM</sub> was able to follow the human path all the way to the target. These observations suggest the necessity of the LSTM layer for the network to acknowledge the existence of the error and correct it when required. Figure 4 (right) shows an example where both networks reached the target, with the LSTM variant imitating the human more closely.

### D. Comparison with Benchmarks

Table I shows the performance of each algorithm we evaluated on the test set. We randomly assigned 2/3 of the data to the training set and 1/3 to the test set. The numbers




Fig. 4: Demonstrations of our network navigating from the start location (green) to the goal location (red). The left figure shows a scenario where  $\text{DeepMoTion}_{Conv}$  fails, while  $\text{DeepMoTion}_{LSTM}$  finds the target. The right figure shows a case where both networks reach the target, with the LSTM variant imitating the human more closely.

TABLE I: Performance Metrics Comparison

	<i>SPD</i>	<i>DTW</i>	<i>Proximity</i>	<i>Collisions</i>	<i>Target</i>
$\text{DeepMoTion}_{LSTM}$	151	39	0.31	0.67	100%
$\text{DeepMoTion}_{Conv}$	732	131	0.25	0.89	69%
SFM [6]	3817	51	0.29	0.26	100%
EMP [17]	15437	1187	0.001	7.69	32%
GRP [8]	334	52	0.18	0.78	84%

in the table are averages over all test examples. The results illustrate the ability of  $\text{DeepMoTion}_{LSTM}$  to imitate humans better than the other benchmark algorithms. *SPD* and *DTW* show that  $\text{DeepMoTion}_{LSTM}$  has the lowest path difference among all the tested algorithms, with the next best algorithm, GRP, showing more than double the path difference.

We note that  $\text{DeepMoTion}_{LSTM}$  and SFM both reach the target on 100% of the trials, while other networks often fail. EMP reaches the target in 32% of trials and GRP reaching the target in 84% of the trials. The *proximity* parameter shows  $\text{DeepMoTion}_{LSTM}$  keeps an average *proximity* of 0.31m to any human, which is safer than the other benchmark algorithms. SFM keeps an average *proximity* of 0.29m despite explicitly weighting its repulsive force to humans higher than the other social forces.

With regard to the number of collisions, SFM has the lowest rate among all the algorithms. This can be explained by the ability of the algorithm to stop in the case of dense crowds, while all the other networks were not trained on any human demonstration that exhibited that type of behavior. We expect  $\text{DeepMoTion}$  to learn to stop and avoid collisions better when trained on more pedestrian data showing a wider set of possible navigation scenarios. We also expect that with enough training data in a variety of environments with different features, e.g., crowd density or crowd speed, the network can learn to perform well even when placed in a new environment that it had not been trained on.

Finally, we note that our network was able to navigate even with a LiDAR range other than the one it was trained

on. All the algorithms above were trained and tested on a LiDAR with a 30m range. To show the ability of our network to generalize to different ranges, we tested its performance with a 6m LiDAR range without any retraining. The network was still able to reach the target on 97% of the trials, with an increase in *DTW* to 47 and a decrease in the number of collisions to 0.51. The decrease in collisions was expected, as the network is observing obstacles in locations that were supposed to be free, and thus the robot becomes more careful.

## VI. CONCLUSION AND FUTURE WORK

We introduced a novel deep imitation learning framework and studied its performance when learning to navigate from human traces. We trained the deep network to predict robot command velocities from raw LiDAR scans without the requirement of any preprocessing or classification of the surrounding objects. Our experiments showed  $\text{DeepMoTion}$ 's ability to generate navigation commands similar to humans and plan a path to the target on all test sets, outperforming all the benchmarks on path difference (*SPD* and *DTW*) and *proximity* metrics, and all except SFM on the number of collisions. In addition, we presented a novel loss function to train the network. The loss function allowed us to accommodate for human motion stochasticity while at the same time enabling the robot to navigate safely. Finally, we presented a comparative assessment that showed the necessity of an LSTM layer for a planning algorithm via a deep neural network, where the robot navigating with the non-LSTM variant of our network was led astray on many test cases.

In the future, we plan to train the network to navigate using raw images instead of LiDAR scans, where we believe the larger bandwidth of the data can help the network understand human motion from their point of view. However, unlike  $\text{DeepMoTion}$ , special consideration has to be taken when training the network with images to provide a navigation model that runs in real time when implemented on a mobile platform.

## REFERENCES

- [1] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese. Social LSTM: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pages 961–971, 2016.
- [2] A. Bera, T. Randhavane, R. Prinja, and D. Manocha. Sociosense: Robot navigation amongst pedestrians with social and psychological constraints. *arXiv:1706.01102*, 2017.
- [3] Y. F. Chen, M. Everett, M. Liu, and J. P. How. Socially aware motion planning with deep reinforcement learning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS*, pages 1343–1350, 2017.
- [4] T. B. Dutra, R. Marques, J. Cavalcante-Neto, C. A. Vidal, and J. Pettr. Gradient-based steering for vision-based crowd simulation algorithms. *Computer Graphics Forum*, 36(2):337–348, 2017.
- [5] T. Fernando, S. Denman, S. Sridharan, and C. Fookes. Soft+ hardwired attention: An LSTM framework for human trajectory prediction and abnormal event detection. *arXiv:1702.05552*, 2017.
- [6] G. Ferrer, A. G. Zulueta, F. H. Cotarelo, and A. Sanfeliu. Robot social-aware navigation framework to accompany people walking side-by-side. *Autonomous Robots*, 41(4):775–793, 2017.
- [7] K. Greff, R. K. Srivastava, J. Koutník, B. R. Steunebrink, and J. Schmidhuber. LSTM: A search space odyssey. *IEEE Transactions on Neural Networks and Learning Systems*, 28(10):2222–2232, 2017.
- [8] E. Groshev, A. Tamar, S. Srivastava, and P. Abbeel. Learning generalized reactive policies using deep neural networks. *arXiv:1708.07280*, 2017.
- [9] D. Helbing and P. Molnar. Social force model for pedestrian dynamics. *Physical Review E*, 51(5):4282, 1995.
- [10] P. Henry, C. Vollmer, B. Ferris, and D. Fox. Learning to navigate through crowded environments. In *Proceedings of the IEEE International Conference on Robotics and Automation, ICRA*, pages 981–986, 2010.
- [11] I. Karamouzas, B. Skinner, and S. J. Guy. Universal power law governing pedestrian interactions. *Physical Review Letters*, 113(23):238701, 2014.
- [12] B. Kim and J. Pineau. Socially adaptive path planning in human environments using inverse reinforcement learning. *International Journal of Social Robotics*, 8(1):51–66, 2016.
- [13] D. V. Lu and W. D. Smart. Towards more efficient navigation for robots and humans. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS*, pages 1707–1713, 2013.
- [14] D. Mehta, G. Ferrer, and E. Olson. Autonomous navigation in dynamic social environments using multi-policy decision making. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS*, pages 1190–1197, 2016.
- [15] U. Patel, E. Hatay, M. D’Arcy, G. Zand, and P. Fazli. Beam: A collaborative autonomous mobile service robot. In *Proceedings of the AAAI Fall Symposium on Artificial Intelligence for Human-Robot Interaction, AI-HRI 2017*, 2017.
- [16] S. Pellegrini, A. Ess, K. Schindler, and L. Van Gool. You’ll never walk alone: Modeling social behavior for multi-target tracking. In *Proceedings of the 12th IEEE International Conference on Computer Vision, ICCV*, pages 261–268, 2009.
- [17] M. Pfeiffer, M. Schaeuble, J. Nieto, R. Siegwart, and C. Cadena. From perception to decision: A data-driven approach to end-to-end motion planning for autonomous ground robots. In *Proceedings of the IEEE International Conference on Robotics and Automation, ICRA*, pages 1527–1533, 2017.
- [18] D. Sankoff and J. Kruskal. *Time Warps, String Edits, and Macromolecules : The Theory and Practice of Sequence Comparison*. Addison-Wesley, 1983.
- [19] E. A. Sisbot, L. F. Marin-Urias, R. Alami, and T. Simeon. A human aware mobile robot motion planner. *IEEE Transactions on Robotics*, 23(5):874–883, 2007.
- [20] R. Triebel, K. Arras, R. Alami, L. Beyer, S. Breuers, R. Chatila, M. Chetouani, D. Cremers, V. Evers, M. Fiore, et al. Spencer: A socially aware service robot for passenger guidance and help in busy airports. In *Proceedings of the Conference on Field and Service Robotics, FSR*, pages 607–622, 2016.
- [21] D. Vasquez, B. Okal, and K. O. Arras. Inverse reinforcement learning algorithms and features for robot navigation in crowds: an experimental comparison. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS*, pages 1341–1346, 2014.
- [22] D. Vasquez, P. Stein, J. Rios-Martinez, A. Escobedo, A. Spalanzani, and C. Laugier. Human aware navigation for assistive robotics. In *Experimental Robotics*, pages 449–462. Springer, 2013.
- [23] M. Veloso, J. Biswas, B. Coltin, S. Rosenthal, T. Kollar, C. Mericli, M. Samadi, S. Brandao, and R. Ventura. Cobots: Collaborative robots servicing multi-floor buildings. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS*, pages 5446–5447, 2012.
- [24] B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey. Maximum entropy inverse reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence, AAAI*, pages 1433–1438, 2008.