

“© 2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.”

Multi-modal Scene-compliant User Intention Estimation in Navigation

Kavindie Katuwandeniya^{1*}, Stefan H. Kiss¹, Lei Shi¹, and Jaime Valls Miro¹

Abstract—A multi-modal framework to generate user intention distributions when operating a mobile vehicle is proposed in this work. The model learns from past observed trajectories and leverages traversability information derived from the visual surroundings to produce a set of future trajectories, suitable to be directly embedded into a perception-action shared control strategy on a mobile agent, or as a safety layer to supervise the prudent operation of the vehicle. We base our solution on a conditional Generative Adversarial Network with Long-Short Term Memory cells to capture trajectory distributions conditioned on past trajectories, further fused with traversability probabilities derived from visual segmentation with a Convolutional Neural Network. The proposed data-driven framework results in a significant reduction in error of the predicted trajectories (versus the ground truth) from comparable strategies in the literature (e.g. Social-GAN) that fail to account for information other than the agent’s past history. Experiments were conducted on a dataset collected with a custom wheelchair model built onto the open-source urban driving simulator CARLA, proving also that the proposed framework can be used with a small, un-annotated dataset.

I. MOTIVATION AND INTRODUCTION

Intention recognition is one of the greatest challenges in intelligent Human Robot Interaction (HRI) [1]. As humans, when interacting with one another, an intuitive understanding of the other agents’ intentions is learnt allowing better collaborative working conditions. When a user and a machine work in a cooperative fashion, the same principles apply so that tighter shared controls can be exercised and instinctively assist each other in achieving the common goal, bypassing the explicit need for extensive and reliable communication between robot and user [2]. Defining user intention is however a challenging proposition. Moreover, it is non-deterministic: given the same situation and user, there could be multiple plausible actions followed through, and they should all be allowed due consideration within a Human-Robot Collaboration (HRC) shared control scheme.

In this work we present a HRC framework for navigational tasks, with the incentive of providing assistance to frail and less able users in driving personal mobility devices (PMDs), such as wheelchairs and scooters. In this context, shared control is perceived differently to other situations where human and robot operate collaboratively on a common workspace, such as assembling pieces together, since actions have a direct impact on the safety and comfort of the rider, and new and unstructured environments are commonplace.

We propose a multi-modal data-driven probabilistic approach where the behaviour of a PMD user captured via the kinematics of the vehicle are combined with contextual




Fig. 1: First person view from the on-board camera depicting the **ground truth** trajectory subsequently followed by the user, alongside the equally probable **predicted user intentions** over the same time horizon, projected onto the image.

awareness provided by the scene captured from an on-board monocular camera to estimate the user intention. A conditional Generative Adversarial Network (GAN) with Long-Short Term Memory (LSTM) encoding and decoding blocks is employed to capture trajectory distributions conditioned on past trajectories. Rejection sampling is then conducted on future trajectory samples from this distribution to fuse scene traversability information derived from visual segmentation with a Convolutional Neural Network (CNN). The final outcome is a set of equally-probable user intention trajectories, graphically illustrated in Fig. 1, suitable for integration into a HRC shared control framework for active vehicles, unlike other alternatives where random selection or an oracle are assumed [3], [4].

It is rare to find an annotated, considerably large dataset to train generative models of one’s domain of work. Our method was developed with the intention of capitalizing on the data available in small, un-annotated datasets.

The paper contributions can be summarised in the development of three components:

- a probabilistic approach for multi-modal user intention estimation (Sec. III),
- a wheelchair implementation in the simulation environment CARLA [5] for safe collection of datasets in realistic urban scenarios (Sec. IV-A), and
- a modular framework able to leverage existing work in the visual segmentation space, suitable to small-scale navigation datasets with no visual annotations (Sec. IV-A and Sec. IV-C).

¹Robotics Institute, University of Technology Sydney, Australia
KavindieHansika.Katuwandeniya@student.uts.edu.au




Fig. 2: Proposed Architecture with the mean of sampled trajectories given in **bold**

II. BACKGROUND

1) *User Intention Estimation*: User intention can be interpreted in numerous ways and is not always straightforward to estimate. Demeester et al. [6] were the first to break free from the standard definitions of user intention in navigation at that time, which in the context of a robotic wheelchairs were mostly specific assistance algorithms such as “follow-corridor”, “avoid-obstacle”, or “pass through door”, to represent them instead as a set of trajectories each with a goal state. In addition, user intention has been reported in the literature as the immediate controls for the mobility device [7], the future trajectory (or control commands of the trajectory) for a predefined time period, or the final goal pose [8]. In this work we consider the intended future trajectory to be the user intention.

2) *Multi-modal Trajectory Prediction*: Trajectories capture the spatial and temporal relationships of an object. Classical methods such as Kalman filters [9] and time-series analysis [10] have been used in the past to estimate them from past and current observations. However, the complex nature of human behaviours impose the need to disregard hand tailored deterministic models as they fall short in capturing the wide range of plausible future trajectories that represent the user intentions. This has given rise to data-driven approaches proven suitable in predicting complex and multi-modal distributions.

Generative models fall under stochastic approaches for data-driven trajectory prediction [11]. Stochastic approaches choose from multiple possibilities based on a random sampling. Given their flexibility and proven ability to model complex distributions, a generative-model-based approach for trajectory prediction was the chosen paradigm implemented in our work.

3) *Generative Models*: Generative models are optimised to capture the underlying data distribution. With the advent of deep learning methods, a new set of generative models were introduced: Deep Generative Models (DGMs), including GANs.

In the context of trajectory prediction, GANs have been

extensively utilised in recent literature. [3] was a pioneer in using GANs in the context of identifying and predicting the behaviour of interacting agents with Social-GAN. [4] uses a GAN with an additional information loss to allow for a disentangled representation between the input to the generator and the sampled output. However, both these methods do not include the scene information in their prediction. More recently, [12] uses a GAN along with scene information to predict the future path of the agents. However, they use a top-down view of the scene which is not possible in our case, or generally for that matter.

Because of instabilities of GANs, predicting trajectories from high dimensional image spaces is difficult. Especially with a small dataset, the chance of over-fitting is high. Thus, it makes sense to use data-driven methods for kinematic-based trajectory prediction and separately perform vision-based scene information integration, where rule-based approaches and pre-trained networks reduce the learning difficulty as they don’t need specialisation for the specific hardware or scenario.

4) *Segmentation networks*: To develop a solution which does not rely on a large, annotated image dataset, and leverage separate advances in this domain, it is advantageous to make use of pre-trained networks. Thus for the rule-based trajectory selection approach, we chose a pre-trained image segmenting network. For a detailed description of the evolution of deep neural network based image-segmentation methods, the reader is referred to [13]. For the purpose of trajectory selection, it is adequate to understand the different classes in a scene. Thus in this work we focus on semantic segmentation techniques trained on a moving vehicle dataset.

III. METHODOLOGY

For this work we define the user intention to be the trajectory the user intends to follow. This definition itself is vague since the trajectory depends on how long it is executed and also the intended destination. Thus it is sensible to define the user intention to be the intended trajectory leading to a point in space at time $t = t + N$ where N is the number of time steps considered into the future. The choice of N

could depend on the Field Of View (FOV) of the user or the safe travel distance or time recommended for the mobile platform.

Mathematically, we denote the position of the wheelchair at discretized time index t by $\mathbf{x}_t = (x_t, y_t)$ where x_t and y_t denote the 2-dimensional position at time t . Given the past M positions $\mathbf{X} = \mathbf{x}_{t+1-M:t}$, we wish to predict the next N positions $\tilde{\mathbf{X}} = \mathbf{x}_{t+1:t+N}$. We also wish to inform our prediction using scene information: an image \mathbf{I} captured at time t from a camera facing forward from the vehicle. Thus, we are interested in modelling the distribution

$$\Pr(\tilde{\mathbf{X}} | \mathbf{X}, \mathbf{I}). \quad (1)$$

We assume that \mathbf{X} and \mathbf{I} are independent, and also conditionally independent given $\tilde{\mathbf{X}}$:

$$\begin{aligned} \mathbf{X} &\perp\!\!\!\perp \mathbf{I}, \\ (\mathbf{X} \perp\!\!\!\perp \mathbf{I}) &| \tilde{\mathbf{X}}. \end{aligned} \quad (2)$$

Given an uninformative uniform prior on $\tilde{\mathbf{X}}$, these assumptions allow us to calculate the posterior distribution as proportional to the product of the two conditional probabilities:

$$\Pr(\tilde{\mathbf{X}} | \mathbf{X}, \mathbf{I}) \propto \Pr(\tilde{\mathbf{X}} | \mathbf{X}) \Pr(\tilde{\mathbf{X}} | \mathbf{I}). \quad (3)$$

A conditional GAN is trained to learn the specific user behaviour via the kinematics of the driven vehicle: $\Pr(\tilde{\mathbf{X}} | \mathbf{X})$. A segmentation network is used to incorporate scene information: $\Pr(\tilde{\mathbf{X}} | \mathbf{I})$. These two sources of information are fused probabilistically to produce a generative model. This approach allows the generative model to be scene compliant when trained with only a small dataset.

A. Trajectory Generative Model

We use a conditional GAN [14] model inspired by [3] to model the user intention from past positions. A GAN framework consists of two adversarial networks: a generator G and a discriminator D , which are jointly optimized such that the generator learns the underlying data distribution, in our case the user intention as a predicted trajectory distribution. The objective function of a GAN can be written as

$$\min_G \max_D V(D, G) = \mathbb{E}_{x \sim \Pr_{data}(x)} [\log D(x)] + \mathbb{E}_{z \sim \Pr_z(z)} [\log(1 - D(G(z)))] \quad (4)$$

where $\Pr_{data}(x)$ is the original data distribution and $\Pr_z(z)$ is a normal distribution from which a noise value is sampled. The discriminator D is fed with either a real data sample x or a generated sample $G(z)$.




Fig. 3: Conditional GAN inspired from Social-GAN [3]




Fig. 4: Scene information. (left): original image. (middle): segmented image, coloured by the most likely class. Legend names the 12 most prominent classes. (right): probability of traversability. White indicates the highest probability of 1.

A conditional GAN is an extended GAN, where the generator and the discriminator are both conditioned on extra information: in our case the observed past trajectory. This framework is shown in Fig. 3. LSTM blocks are used to capture the temporal relationships of positions. The objective function of the GAN is modified as

$$\min_G \max_D V(D, G) = \mathbb{E}_{x \sim \Pr_{data}(x)} [\log D(x|y)] + \mathbb{E}_{z \sim \Pr_z(z)} [\log(1 - D(G(z|y)))] \quad (5)$$

where y conditions both networks G and D .

The input for our GAN is the past observed trajectory $\mathbf{X} = \mathbf{x}_{t+1-M:t}$ in the form of consecutive relative positions. In addition to this, a noise vector \mathbf{z} is independently sampled from a normal distribution and fed as an input. The output of the generator is a future trajectory in the form of consecutive relative positions, $\tilde{\mathbf{X}} = \mathbf{x}_{t+1:t+N}$. Different noise inputs result in different predictions, leading to a distribution of possible output trajectories. Resulting trajectories of this network are shown in blue in Fig. 2 where $k = 20$ samples were generated.

The data distribution the GAN framework optimizes to learn is $\Pr(\tilde{\mathbf{X}} | \mathbf{X})$ and with different noise values, it outputs different samples from this learnt distribution,

$$\tilde{\mathbf{X}} = G(\mathbf{z}, \mathbf{X}) \sim \Pr(\tilde{\mathbf{X}} | \mathbf{X}). \quad (6)$$

The main drawback in this framework is that it does not take into account the scene understanding.

B. Incorporating Scene Information

We integrate the visual information by segmenting the image, generating a probability distribution over pixel-wise classes \mathbf{C} . Gated-SCNN [15] was used as the image segmentation network. Further details of the implementation is given under Sec. IV-C. Fig. 4 shows the original and segmented images as the left and middle images respectively.

The segmentation network S estimates the conditional probability of the scene pixel classes \mathbf{C} given the image \mathbf{I} ,

$$\mathbf{C} \sim S(\mathbf{I}) = \Pr(\mathbf{C} | \mathbf{I}). \quad (7)$$

The N waypoints of a predicted trajectory $\tilde{\mathbf{X}}$ are mapped onto the image plane via projective transform $H(\cdot)$. Vehicle positions are assumed to be 2-dimensional lying on the

ground plane ($z = 0$) as motion is safely assumed to happen over relatively flat terrain. The image coordinate \mathbf{u} of trajectory position \mathbf{x} is given by

$$\begin{bmatrix} \mathbf{u} \\ 1 \end{bmatrix} = \mathbf{K} \begin{bmatrix} \mathbf{R} & \mathbf{t} \\ \mathbf{0} & 1 \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{x} \\ 0 \\ 1 \end{bmatrix} \quad (8)$$

$$\mathbf{u} = H(\mathbf{x}),$$

where \mathbf{K} is the intrinsic camera matrix and \mathbf{R} and \mathbf{t} correspond respectively to the 3-dimensional rotation and translation of the camera in world-coordinates.

Out of the 19 class labels that the Gated-SCNN was trained on (more details in Sec. IV-C), “road” and “sidewalk” were considered traversable and all other class were categorized as non-traversable. Based on this we can assign a probability of traversability to each pixel. An example is shown as the *right* image of Fig.4. By projecting a position \mathbf{x} into the image frame, this constraint can be enforced as the conditional probability distribution

$$\Pr(\mathbf{x} | \mathbf{C}) \propto \begin{cases} 1, & \text{if } \mathbf{C}(H(\mathbf{x})) \in \mathbf{C}_{\text{Traversable}} \\ 1, & \text{if } \mathbf{x} \text{ is at the foot of the robot} \\ 0.5, & \text{if } \mathbf{x} \text{ is outside the visible area} \\ 0, & \text{otherwise} \end{cases} \quad (9)$$

where $\mathbf{C}_{\text{Traversable}}$ is the set of traversable segmentation classes $\{\mathbf{C}_{\text{road}}, \mathbf{C}_{\text{sidewalk}}\}$. The notation $\mathbf{C}(H(\mathbf{x}))$ gives the segmentation class for the pixel at position \mathbf{x} projected into the image frame.

Pixels in the segmented image are assumed to be independent, therefore the joint probability of the future trajectory $\tilde{\mathbf{X}}$ is computed as the product

$$\Pr(\tilde{\mathbf{X}} | \mathbf{C}) \propto \prod_{\mathbf{x} \in \tilde{\mathbf{X}}} \Pr(\mathbf{x} | \mathbf{C}) \quad (10)$$

Thus, $\Pr(\tilde{\mathbf{X}} | \mathbf{I})$ can now be found by,

$$\Pr(\tilde{\mathbf{X}} | \mathbf{I}) = \sum_{\mathbf{C}} \Pr(\tilde{\mathbf{X}} | \mathbf{C}) \Pr(\mathbf{C} | \mathbf{I}). \quad (11)$$

C. Probabilistic Fusion

We fuse the information about the future trajectory $\tilde{\mathbf{X}}$ from the past trajectory \mathbf{X} and the camera image \mathbf{I} to refine Eq. (3),

$$\Pr(\tilde{\mathbf{X}} | \mathbf{X}, \mathbf{I}) \propto \underbrace{\Pr(\tilde{\mathbf{X}} | \mathbf{X})}_{G(\mathbf{z}, \mathbf{X})} \sum_{\mathbf{C}} \prod_{\mathbf{x} \in \tilde{\mathbf{X}}} \underbrace{\Pr(\mathbf{x} | \mathbf{C})}_{\text{traversability constraint}} \underbrace{\Pr(\mathbf{C} | \mathbf{I})}_{S(\mathbf{I})} \quad (12)$$

Rejection sampling is used to form a generative model from the posterior distribution $\Pr(\tilde{\mathbf{X}} | \mathbf{X}, \mathbf{I})$. Treated as a proposal distribution, the conditional distribution $\Pr(\tilde{\mathbf{X}} | \mathbf{X})$ is sampled from the GAN conditioned on the past trajectory $G(\mathbf{z}, \mathbf{X})$. The second conditional probability $\Pr(\tilde{\mathbf{X}} | \mathbf{I})$ is evaluated using the segmentation of the camera image $S(\mathbf{I})$ and considering the traversability constraint $\Pr(\tilde{\mathbf{X}} | \mathbf{C})$. The sample is accepted proportional to $\Pr(\tilde{\mathbf{X}} | \mathbf{I})$. This procedure is repeated until the required number of accepted samples are generated. As rejection sampling is scale invariant, the

proportionality relationships of previous equations are sufficient for evaluation. The overall framework is illustrated in Fig. 2 with the proposed trajectories shown in red.

IV. DATASET AND IMPLEMENTATION

A. CARLA Dataset

The dataset used for training the GAN framework (Sec. IV-B) was generated within CARLA [5], an open-source simulation environment for realistic and safe testing of vehicles driving in urban layouts, with humans and vehicles on the road following natural behaviours. A personal mobility device was created and added to the environment for the purpose of this work. A Robotic Operating System (ROS) [16] bridge was used to interface CARLA with ROS and an external joystick was used to control the wheelchair. The wheelchair was driven around the environment collecting data for training.

The dataset consists of six ROS bags each with an average duration of 10 minutes. One bag was kept for validating and one for testing, the remaining four were used for training. The training data was organised in 0.5s sampling intervals which resulted in 5252 training instances, 1308 validating instances and 1261 testing instances. Even though 5252 training instance were selected, the environment itself is fairly homogeneous, lacking variation. The wheelchair was driven around city blocks resulting in a relatively small area which produced similar urban imagery in the dataset. This lack of diversity in the training instances is believed to be the culprit for the over-fitting we experienced when end-to-end GANs with images as sole input were first tested. Moreover, the data is not annotated and manual labelling would have been a significant time-consuming exercise.

B. GAN Network Implementation

The GAN framework was inspired by Social-GAN [3] as in Fig. 3. The integration of LSTM cells allows the network to learn temporal relationships. The generator network consists of an encoder and a decoder network. The encoder encodes the past observed trajectory \mathbf{X} as consecutive relative positions as an encoded vector to which an 8 dimensional noise vector is added. For this work, M was selected to be 8 with a 0.5s period between time steps. Thus, \mathbf{X} consists of 4s of past vehicle trajectory up to the current time step. Each dimension of the noise vector is randomly sampled from a normal distribution independently.

The hidden layer of the decoder is initialized with the encoded \mathbf{X} and the noise vector. The LSTM cell outputs the consecutive relative future positions as the output from the generator network: $\tilde{\mathbf{X}}$. As with \mathbf{X} , N is selected to be 8, such that the generator will output a trajectory 4s in length starting from the immediate next time step.

The discriminator takes as input the past trajectory \mathbf{X} with either the ground truth or the predicted future trajectory $\tilde{\mathbf{X}}$, and sends them through LSTM cells to output a scalar value which is then compared against 1 or 0: real or fake. As an implementation detail, the labels are smoothed such that $fake \sim \text{Uniform}(0, 0.3)$ and $real \sim \text{Uniform}(0.7, 1)$. D is

	Selection	No scene information $\Pr(\tilde{\mathbf{X}} \mathbf{X})$	Proposed $\Pr(\tilde{\mathbf{X}} \mathbf{X}, \mathbf{I})$	Improvement
ADE (m)	random	1.8844	1.6094	-14.59%
	mean	1.6020	1.3617	-15.00%
	\min_k	0.5406	0.5079	-6.05%
FDE (m)	random	4.0259	3.3507	-16.77%
	mean	3.3489	2.7988	-16.43%
	\min_k	0.7416	0.6700	-9.65%

TABLE I: Prediction errors with respect to ground truth.

trained with the binary cross entropy loss of the generated samples and fake labels; and ground truth trajectories and real labels. G is trained using the binary cross entropy loss between the generated samples and real labels. In addition to this variety loss is also considered, where the minimum L2 loss of the generated k samples and the ground truth trajectory is also back-propagated. The Adam optimizer was used to train both networks. The network setup and training procedure followed the Social-GAN framework [3]. The implementation was done in PyTorch [17].

C. Segmentation Network Implementation

Gated-SCNN [15] was used as the image segmentation network. This model is trainable end-to-end and is trained on the urban Cityscapes dataset [18] which has 34 semantic labels (of which 19 were used during the training process) from images captured from a road-driving vehicle. Gated-SCNN is a state-of-the-art model, outperforming similar semantic segmentation models considering mean average precision (mAP) as Intersection over Union (IoU) threshold [13]. Moreover, the implementation and trained network weights are publicly available and implemented in PyTorch. However, any other semantic segmentation model can be used to replace the Gated-SCNN and the accuracy will change as per the selected model.

The model consists of 2 streams of networks: a regular stream and a shape stream. The regular stream processes semantic region information using a classical CNN. The shape stream uses low-level feature maps from the regular stream to process the boundary information using Gated Convolution Layers (GCLs). The output of both streams are fed to a fusion module, which outputs semantic regions with clear boundaries. No training of the Gated-SCNN was conducted, the pre-trained model weights were used to semantically segment the image \mathbf{I} at time t .

V. RESULTS AND ANALYSIS

Average Displacement Error (ADE) and Final Displacement Error (FDE) are common metrics for uni-modal trajectory prediction; ADE gives the mean Euclidean distance along a fixed-duration trajectory while FDE gives the distance error associated to the trajectory endpoints. According to [19], there is no agreed upon metric when it comes to comparing multi-modal prediction frameworks. Table I shows the improvement of the proposed method with scene understanding to the original (no scene information) Social-GAN framework with different trajectory selection strategies.




Fig. 5: Analysis of the \min_k metric. ADE (left) and FDE (right) of the trajectory closest to the ground truth selected from k samples. Error with the original set-up ($\Pr(\tilde{\mathbf{X}}|\mathbf{X})$) and error with the proposed framework ($\Pr(\tilde{\mathbf{X}}|\mathbf{X}, \mathbf{I})$) shows that our improvement is consistent across the choice of k .

As with previous work [19], [3], [4], we shall use ADE \min_k and FDE \min_k as evaluation metrics. Out of the k samples, the sample which results in the lowest ADE and FDE are used when comparing the results. With the proposed method, there is a 6.05% decrease in the ADE \min_k error and 9.65% decrease in the FDE \min_k metric as in Table I. A comparison figure depicting how the ADE \min_k and FDE \min_k changes with the number of samples selected is plotted in Fig. 5, reflecting how the decreasing in the error with the proposed method is consistent across any choice of k samples.

In a real world deployment, if the system is to estimate the user intention and provide the required control assistance based on the estimated intention, one trajectory from the generated k has to be selected. Picking the trajectory closest to the ground truth becomes infeasible. One could either select a random sample (equivalent to \min_1) or pick the mean of all sampled trajectories. Our framework shows a consistent decrease in both ADE and FDE under either of these scenarios as per Table I under *random* and *mean* selection respectively.

Two instances illustrating typical navigation scenarios are shown in Fig. 6. If the system was to consider only the kinematics of the assistive device, it would generate the **blue** trajectories (hereafter the mean of a trajectory distribution will be depicted in **bold** when projected onto the images). Using the proposed method with consideration for the image information results in (the same number of) equally probable **red** trajectories, compliant with the traversability insights embedded in the image.

Fig. 6 qualitatively hints at the potential advantage of incorporating visual information into a trajectory generative framework. Fig. 6a shows that by rejecting forward- and right-directed trajectories that would have collided with the seat or building, the model predicts more plausible left-directed trajectories, closer to the ground truth. Fig. 6b similarly shows a strong increase in left-directed trajectories, but also highlights a weakness of the proposed framework. It can be seen in the projected image that the ground-truth trajectory passes behind a pole, which the segmentation classified as non-traversable. The flawed assumption is that all pixels in the segmented image lie on the ground plane. Various image processing techniques or manipulations to the traversability constraint $\Pr(\tilde{\mathbf{X}}|\mathbf{C})$ are likely able to correct

ACKNOWLEDGEMENT

The authors would like to thank Janmart Tenedora for his contribution in developing the CARLA simulation environment and the collection of the dataset.

REFERENCES

- [1] J. L. Burke, R. R. Murphy, E. Rogers, V. J. Lumelsky, and J. Scholtz, "Final report for the darpa/nsf interdisciplinary study on human-robot interaction," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 34, no. 2, pp. 103–112, 2004.
- [2] K. A. Tahboub, "Intelligent human-machine interaction based on dynamic bayesian networks probabilistic intention recognition," *Journal of Intelligent and Robotic Systems*, vol. 45, no. 1, pp. 31–52, 2006.
- [3] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, "Social GAN: Socially Acceptable Trajectories with Generative Adversarial Networks," *Proc. of the IEEE Computer Society Conference on CVPR*, pp. 2255–2264, 2018.
- [4] J. Amirian, J. B. Hayet, and J. Pettre, "Social ways: Learning multimodal distributions of pedestrian trajectories with GANs," *IEEE Computer Society Conference on CVPR Workshops*, vol. 2019-June, pp. 2964–2972, 2019.
- [5] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," in *Conference on robot learning*. PMLR, 2017, pp. 1–16.
- [6] E. Demeester, A. Huntemann, D. Vanhooydonck, G. Vanacker, A. Degeest, H. Van Brussel, and M. Nuttin, "Bayesian estimation of wheelchair driver intents: Modeling intents as geometric paths tracked by the driver," in *2006 IEEE/RSJ IROS*. IEEE, 2006, pp. 5775–5780.
- [7] K. Katuwandeniya, J. V. Miro, and L. Dantanarayana, "End-to-end joint intention estimation for shared control personal mobility navigation," in *2020 16th ICARCV*. IEEE, 2020, pp. 1–6.
- [8] V. K. Narayanan, A. Spalanzani, and M. Babel, "A semi-autonomous framework for human-aware and user intention driven wheelchair mobility assistance," in *2016 IEEE/RSJ IROS*. IEEE, 2016, pp. 4700–4707.
- [9] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 03 1960.
- [10] M. B. Priestley, *Spectral analysis and time series: probability and mathematical statistics*. Academic Press, 1981, no. 04; QA280, P7.
- [11] T. Phan-Minh, E. C. Grigore, F. A. Boulton, O. Beijbom, and E. M. Wolff, "Covernet: Multimodal behavior prediction using trajectory sets," in *Proc. of the IEEE/CVF Conference on CVPR*, 2020, pp. 14 074–14 083.
- [12] J. Li, H. Ma, and M. Tomizuka, "Conditional Generative Neural System for Probabilistic Trajectory Prediction," in *IEEE IROS*, 2019, pp. 6150–6156.
- [13] F. Sultana, A. Sufian, and P. Dutta, "Evolution of image segmentation using deep convolutional neural network: a survey," *Knowledge-Based Systems*, vol. 201, p. 106062, 2020.
- [14] M. Mirza and S. Osindero, "Conditional generative adversarial nets," *arXiv preprint arXiv:1411.1784*, 2014.
- [15] T. Takikawa, D. Acuna, V. Jampani, and S. Fidler, "Gated-scnn: Gated shape cnns for semantic segmentation," in *Proc. of the IEEE/CVF ICCV*, 2019, pp. 5229–5238.
- [16] Stanford Artificial Intelligence Laboratory et al., "Robotic operating system." [Online]. Available: <https://www.ros.org>
- [17] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in *NeurIPS*. Curran Associates, Inc., 2019, pp. 8024–8035.
- [18] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The cityscapes dataset for semantic urban scene understanding," in *Proc. of the IEEE conference on CVPR*, 2016, pp. 3213–3223.
- [19] J. Liang, L. Jiang, K. Murphy, T. Yu, and A. Hauptmann, "The garden of forking paths: Towards multi-future trajectory prediction," in *Proc. of the IEEE/CVF Conference on CVPR*, 2020, pp. 10 508–10 518.
- [20] L. Shi, S. Kodagoda, and M. Piccardi, "Towards simultaneous place classification and object detection based on conditional random field with multiple cues," in *2013 IEEE/RSJ IROS*. IEEE, 2013, pp. 2806–2811.

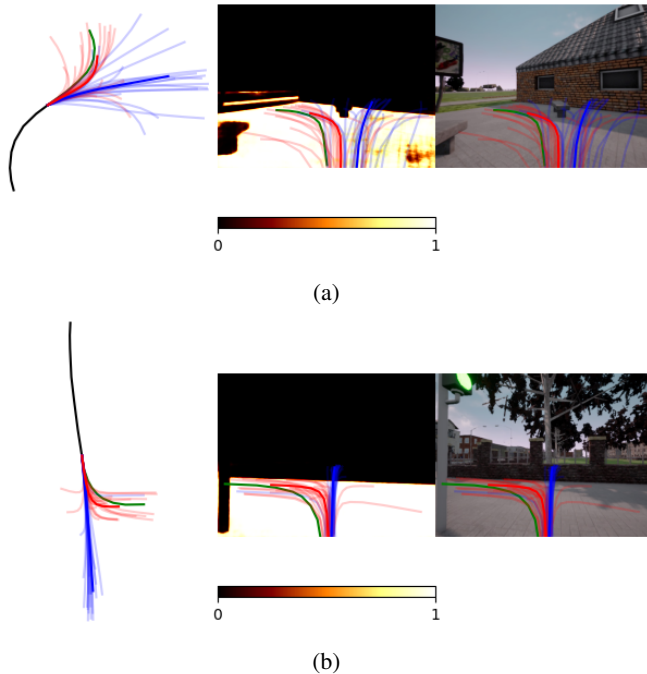


Fig. 6: Two examples illustrating the proposed methodology. (left) all the trajectories in the world coordinate frame where **observed past trajectory**, **ground truth trajectory**, **predicted and accepted trajectories (mean)**, **predicted but rejected trajectories (mean)**, (middle) the same trajectories projected onto the segmented image plane. (right) all the trajectories projected onto the real image plane.

this shortcoming, and this work is left for future work.

In this work we have assumed independence for traversability, relying on the traversability generated from the current image to produce $\Pr(\tilde{\mathbf{X}} | \mathbf{I})$. The framework can be improved further by incorporating spatial dependence along the generated trajectory, and utilising more environmental features, which could be learned from historically labelled ground-truth trajectories, e.g. with Conditional Random Fields (CRFs) [20].

VI. CONCLUSION

A framework is developed to estimate user intentions in the space of HRC between a mobility robotic agent and its user. A crucial aspect of this setup is the possibility of having multiple equally-probable estimates for the future trajectory. The framework considers the user behaviour through the kinematics of the vehicle and leverages the scene context separately, probabilistically fusing both information sources to generate the final trajectory distribution. This allows leveraging on separate advances in the image segmentation domain, whilst comparing favourably in its ability to generate trajectories from a small-scale navigation dataset against comparable data-driven schemes in the literature, generating error improvements in the order of 6% or above.