Interactive Pedestrian Simulation in iGibson

Yige Liu*, Siyun Li*, Chengshu Li, Claudia Pérez-D'Arpino, Silvio Savarese
Stanford University

***Abstract*—We present a learning-based local controller for pedestrian simulation in iGibson. We explore the feasibility of leveraging datasets of pedestrian trajectories to learn a model for socially-aware pedestrian simulation in indoor constrained environments. Based on previous work on pedestrian prediction, we augment Social-GAN with information on static obstacles, pedestrian trajectory histories, and pedestrian goals in order to use the architecture for human-like trajectory generation. In ongoing work, we empirically find that his local controller exhibits more realistic interactive behaviors than commonly used models in social navigation research such as ORCA.**

I. Introduction

Recent work on robot learning for social navigation leverages simulation environments within realistic indoor spaces populated with simulated pedestrians [16], such as the simulation shown in Fig.1. The protocol for optimal reciprocal collision avoidance (ORCA), which is frequently used in social navigation studies [5], generates pedestrian motion commands per time step by solving a linear program with known velocities of all the agents, and guarantees multi-agent collision avoidance. Each agent computes its velocity independently with no explicit communication (other than full observability of all instant velocities). However, the ORCA protocol is limited for simulating pedestrians in indoor constrained environments with concave obstacles, turns or complex high-level planning. ORCA requires coordination among all agents [16], dividing the responsibility for collision avoidance equally among the interacting agents, with no explicit modelling of human-like social interactions. ORCA also fails to address the multi-modal nature of pedestrian behaviors.

We focus on developing a learning-based socially-aware local controller for pedestrian simulation. By incorporating information on static obstacles, pedestrian trajectory histories, and pedestrian goals, we empirically find that this local controller exhibits more realistic interactive behaviors. Our method uses a pedestrian prediction method as a generative model for interactive pedestrian trajectories, by augmenting social-GAN [8] with additional inputs for pedestrians goals and environment map, resulting on a method that is goal oriented and considers static obstacles and the physical constraints of the layouts. Our proposed architecture is shown in Fig. 2.

II. Related Work

Bringing dynamic pedestrians into simulated robotic training environments has been identified to be crucial for robots to learn socially complaint navigation policies. Related work covers both analytical and learning based methods.

Fig. 1: Top view of a simulated environment in iGibson, with two pedestrians moving towards a door. This common situation in indoor environments requires the simulated pedestrians to exhibit human-like behaviors.

A. Analytical Simulation

There are two major approaches to modeling pedestrian dynamics in existing simulators. The first approach is the social force model proposed by Helbing and Molnar [10], which has been implemented in simulators such as Gazebo [19]. The social forces model is based on the physics rule that the accelerations of the pedestrians are proportional to the sum of forces applied to them. This set of forces is modelled after social behaviors. Since the social force model is based on handcrafted features, it could only reflect the pre-defined set of pedestrian navigation rules.

There are also significant efforts in constructing complete pipelines for pedestrian simulation including high level planning, such as in Menge [6]. Similar to our approach, Menge decouples the pedestrian simulation pipeline into four modules: goal selection, plan computation, plan adaptation, and motion synthesis. We seek a learning method with the goal of achieving simulated behaviors that resemble those of human navigation.

B. Learning-Based Trajectory Forecasting

Pedestrian trajectory forecasting algorithms leverage learning methods including Generative Models [8][19][2][17], LSTM (encoder-decoder) based algorithms [1][3], and variational learning approaches [4] to forecast socially acceptable human trajectories with the generator or decoder. A significant

[0]* indicates these authors contributed equally.
RSS 2021 Workshop on Social Robot Navigation.

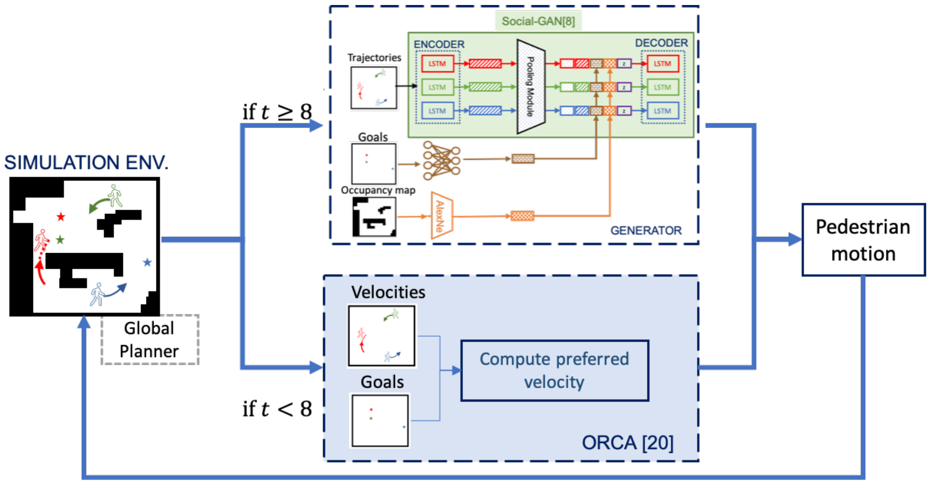

Fig. 2: Overview of our pedestrian simulation pipeline. The iGibson global planner creates a set of waypoints that the local planner uses as short-term goals. The local planner also takes the occupancy map of iGibson's indoor environment and trajectory histories of each pedestrians as input and predicts the next positions for each pedestrian. The trajectory history is initialized by running the default ORCA simulator for first 7 time steps (i.e. $t < 8$) of each episodes.

advantage of the learning method is its ability to generate pedestrian behaviors based on real human data and avoid the need for hand-crafting properties of pedestrians. In addition, learning approaches address the multi-modality nature [4] of human trajectories.

Existing methods of pedestrian trajectory forecasting mainly focus on modeling the social behaviors and they do not have access to privileged scene information that is only pertinent to simulation environments. To this end, we adapt the pedestrian forecasting algorithm to the task of pedestrian simulation for iGibson [18] with the additional knowledge of (1) the static environment surrounding the pedestrians, (2) the navigation goals of the pedestrian, and (3) the desired velocity offered by the global planner. This allows our local path planner to navigate the pedestrians without collisions in iGibson-specific indoor environments while still maintaining social interactions between pedestrians.

III. Methods

We follow a hierarchical paradigm for simulating interactions between pedestrians in iGibson indoor environments. In particular, the hierarchy is composed of two levels: *Global Planning* to construct coarse-grained waypoints to guide pedestrians towards their goals, and *Local Planning* to refine global waypoints and control pedestrian positions exactly at each timestep. Such pipeline, combining global plan generator with learning models for local-replan, is also found to be effective in robot navigation problems for generating socially complaint actions towards the goal [15]. Figure 2 shows an overview of our pedestrian simulation pipeline.

A. Global Planning for Waypoint Generation

The global planner constructs a graph that connects all open locations on the occupancy map with weighted edges equal to the L_2 distance between each pair of nodes.

For each pedestrian i, we then sample a goal position $p_{i,goal} = (x_{i,goal}, y_{i,goal})$ and a start position $p_{i,0} = (x_{i,0}, y_{i,0})$ such that $p_{i,goal}, p_{i,0} \in S_{open}$. The global planner takes in these two positions and runs A^* algorithm [9] to find the shortest path to reach goal.

B. Local Planning for Social Trajectory Generation

Socially acceptable behavior is difficult to model exactly using analytical methods. To this end, we leverage learning-based generative methods to imitate human trajectories collected from real-world data. Figure 3 shows the architecture for our local planner.

1) Problem Definition: Our goal is to simultaneously predict the positions of all n pedestrians in the next time step. The local planner takes as input the trajectory history for all pedestrians in the scene as $X = X_1, X_2, ...X_n$, a floorplan I as the occupancy map, and a set of waypoints leading to $p_{goal} = p_{1,goal}, p_{2,goal}, ...p_{n,goal}$. The local planner will then predict the future positions as $\hat{Y} = Y_1, Y_2, ...Y_n$. For each pedestrian i, its trajectory history at time $t \geq 8$ is defined as $X_i = \{(x_i^{(t-7)}, y_i^{(t-7)}...(x_i^t, y_i^t)\}$ and its goal position is given by the global planner. The trajectory history for $t < 8$ is collected by running the default ORCA simulator. The pedestrian's position prediction for time $t+1$ is denoted by $\hat{Y}_i = (x_i^{(t+1)}, y_i^{(t+1)})$ and during training, its ground-truth next position is denoted by Y_i.

2) Learning Social Interactions using GAN: Inspired by Social GAN [8], our local planner uses the same social pooling module and the GAN-based encoder-decoder architecture. The generator takes as input the trajectory history X, static occupancy map I and short-term goal positions g_{goal} and outputs predicted future positions $\hat{Y}$. In both the generator and discriminator, each pedestrian trajectory history is encoded by a separate sequence of LSTM cells. To encode the social interactions between pedestrians, we use the same social pooling module proposed by Social GAN. The extracted social feature is then concatenated with the output from the last LSTM cell in the encoder.

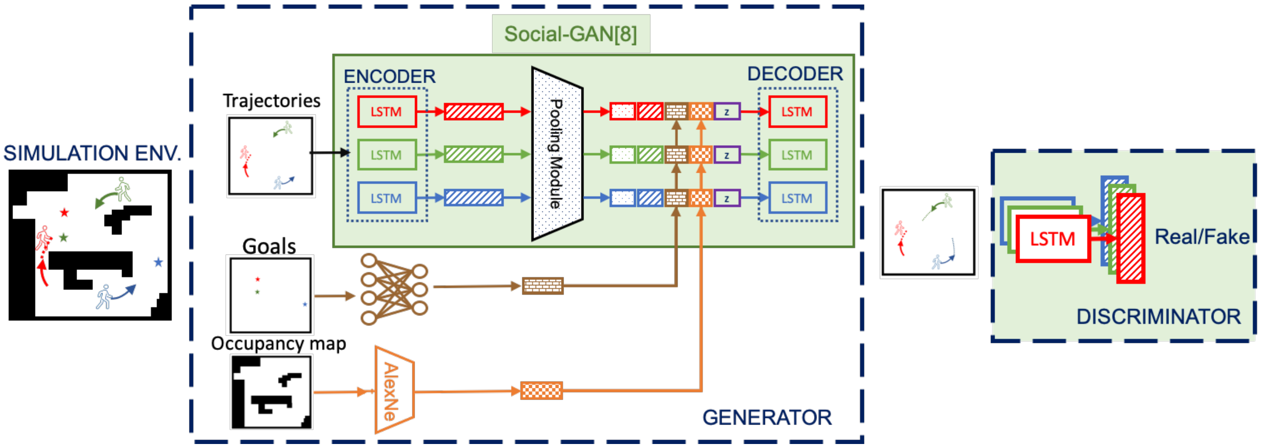

Fig. 3: Proposed model. We use the GAN-based encoder-decoder structure and pooling layer from Social GAN [8] as our backbone model, and incorporate the knowledge of pedestrian short-term goal positions and local obstacles so that the local controller can better follow the global plan while avoiding local collisions.

The decoder generates the future trajectories conditioned on the concatenated features. During training, the discriminator takes two sequences: X with $\hat{Y}$ and X with ground-truth Y, and classifies them as real or synthetic (unlikely to come from real dataset). This encourages the generator to output future trajectory predictions that are consistent with the ground-truth human trajectories, i.e. imitate the social behavior in human navigation.

3) Local Navigation Using Immediate Goals: Unlike traditional trajectory forecasting datasets, our pedestrians have pre-defined goals that they are trying to reach. Since the global planner has provided a list of coarse waypoints, we define the short-term goal positions of each pedestrian as g_i. The coarse waypoints are achieved through joining collinear dense waypoints based on the shortest path to the final goal computed with the traverse map of the environment. For each pedestrian, their goal position is projected to a higher dimension space through a linear layer and the learned representation is concatenated with the LSTM outputs and the noise term as the input to decoder, inspired by [7] and [11].

Note that different from [11] that predicts, the intermediate goals are known in the simulation environment; therefore, we replaced the probability distribution of pedestrian's intent in [11] with global planner outputs.

4) Extracting Obstacle Information using CNN: One challenge for simulating and navigating pedestrians in iGibson is that the indoor environment is cluttered with static objects. Instead of using ORCA [20] to avoid collisions, we choose to use a learning-based perception system to achieve the same goal. Specifically, we use a pre-trained AlexNet [12] model to extract obstacle information from the occupancy map and concatenate the learned map feature with the LSTM outputs before the decoder.

IV. Experiments and Results

In this section, we will describe three sets of experiments and results to evaluate this learning-based socially-aware method.

A. Experiment 1: Benchmark with Trajectory Forecasting

Since we base our method on a trajectory forecasting method, we first evaluate the proposed architecture in a pedestrian prediction task and validate its functionality. We adapt the Social GAN architecture and train our model using ETH [14] and UCY [13] datasets. We use standard trajectory forecasting benchmarks to evaluate our model performance. Specifically, we use Average Displacement Error (ADE) and Final Displacement Error (FDE). We compare our results against two baselines: 1) Social GAN, since our model is a variant of it; and 2) Social-VRNN since it also uses a global occupancy map for producing predictions. Table I shows the results of leave-one-out (dataset) validation. We train the models on all ETH/UCY datasets except for the one we are evaluating on (ZARA01), following the same conventions as that of [8]. In addition, we obtain the quantitative results with 8 time steps of history trajectory observations ($t_{obs} = 8$), and predict 8 steps of future positions ($t_{pred} = 8$). We find that the proposed architecture has comparable or improved results over other prediction methods.

B. Experiment 2: Reaching Goals in iGibson

Although our problem is similar to trajectory forecasting setup, our setting is different in several ways: 1) iGibson is a simulation environment with indoor scenes whereas ETH/UCY recorded in scenes; 2) the pedestrian trajectory forecasting model trained on ETH/UCY datasets is associated to fixed length pedestrian trajectories while in iGibson, the episode length is variable and our pedestrians have a goal position to reach.

To evaluate whether our local controller can move pedestrians toward their goals, we conduct several trials of the experiment in iGibson with the different number of pedestrians spawned in the scene and count the number of times that the pedestrians reach their pre-defined goals. Results in Table II indicate that all of the pedestrians in all of the scenarios could achieve their goals in a reasonable amount of time comparing to the default setting of social navigation task, which is 100

Metric	SGAN	Social-VRNN	SGAN + goal	SGAN + goal + image (partial)
FDE	0.42	0.70	0.07	0.27
ADE	0.21	0.41	0.07	0.17

TABLE I: Pedestrian prediction results of all methods. The methods are trained on ETH/UCY trajectory forecasting benchmark and evaluated on ZARA01 dataset. We report FDE (Final Displacement Error) and ADE (Average Displacement Error) metrics for $t_{pred} = 8$. Our approaches consistently outperform the baseline models in both of the metrics (lower is better).

Num Pedestrians / Num Goals	Total Time Steps	Parity (ped id : num goals reached)
2/3	227	0: 3, 1: 3
2/5	300	0: 5, 1: 8
3/3	139	0: 3, 1: 4, 2: 3
3/5	280	0: 5, 1: 5, 2: 5
5/3	265	0: 5, 1: 5, 2: 5, 3: 5, 4: 7

TABLE II: Quantitative results measuring the pedestrians' abilities of reaching their final goals in iGibson environment with the modified SGAN model and the simulation pipeline. Pedestrians in all of the scenarios reach their goals in reasonable time spans. The parity measurement shows that no pedestrian is blocked from reaching the goals with our simulation approach.

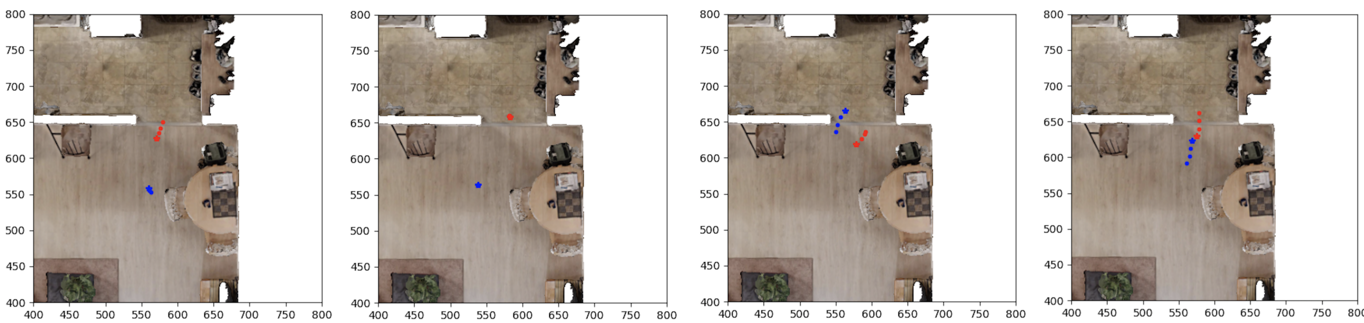

(a) Modified Social GAN (b) $ORCA\_radius = 0.5m$ (c) $ORCA\_radius = 0.1m$ (d) $ORCA\_radius = 0m$

Fig. 4: Sample trajectory plots of two pedestrians going across the door from different directions in iGibson environment. "*" denotes the most recent time step and "**.**" denotes three history time steps. Only modified SGAN approach exhibit a collision-free and socially acceptable manner in the scenario.

time steps per episode

C. Experiment 3: Evaluating Social Interaction

The main objective for our local planner is to add social interaction between pedestrians in simulation. While analytical methods like ORCA are effective in avoiding local obstacles, they do not take social behaviors into consideration. In this experiment, we compare the paths two pedestrians take when going through a door generated by ORCA and our model.

We simulate the same scenario with 4 different pedestrian controllers/controller settings, shown in Figure 4b. The scenario consists of two pedestrians with predefined initial positions and goals that drive the pedestrians to navigate to the other side of the door from different initial directions.

As shown in Figure 4b, when the two pedestrians are controlled by the baseline ORCA algorithm using default ORCA radius of 0.5m, the pedestrians are blocked due to the lack of free space for ORCA's collision avoidance mechanism (across the narrow door as shown in Figure 1. In addition, if we decrease the ORCA radius to 0.1m, although the two pedestrians could successfully reach their goals, as shown in Figure 4c, the pedestrians' trajectories form a circular pattern with a radius around 0.1m. This is because ORCA distributes the collision avoidance responsibilities equally among the pedestrians.

If we use the same simulation scenario but replace ORCA with our learning-based local controller, we can see that the blue pedestrian stops to let the red pedestrian with faster speed to pass first. Moreover, the pedestrians successfully figure out collision-free and socially acceptable movements to pass through the door without sticking their trajectories to a circular pattern.

V. Conclusion

We follow a hierarchical paradigm for pedestrian path planning and control in iGibson. The global planner generates a list of waypoints that guide each pedestrian towards their goal while the learning-based local planner refines the local plan of each pedestrian. The results indicate that even when training on ETH/UCY dataset collected in open space, our local planner can learn from real human trajectory data and still shows human-like and socially acceptable interactions in indoor navigation tasks. Moreover, it shows improvement over ORCA in the socially-relevant test case of two pedestrians going through the door. Videos of the simulation are available at https://sites.google.com/view/int-pedestrians/.

REFERENCES

[1] Alexandre Alahi, Kratarth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social LSTM: Human trajectory prediction in crowded spaces. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 961–971, 2016.

[2] Javad Amirian, Jean-Bernard Hayet, and Julien Pettré. Social ways: Learning multi-modal distributions of pedestrian trajectories with gans. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2964–2972, 2019.

[3] Stefan Becker, Ronny Hug, Wolfgang Hübner, and Michael Arens. *RED: A Simple but Effective Baseline Predictor for the TrajNet Benchmark: Munich, Germany, September 8-14, 2018, Proceedings, Part III*, pages 138–153. 01 2019.

[4] Bruno Brito, Hai Zhu, Wei Pan, and Javier Alonso-Mora. Social-vrnn: One-shot multi-modal trajectory prediction for interacting pedestrians, CORL 2020.

[5] Yuhang Che, Allison M. Okamura, and Dorsa Sadigh. Efficient and trustworthy social navigation via explicit and implicit robot–human communication. *IEEE Transactions on Robotics*, 36(3):692–707, 2020. doi: 10.1109/TRO.2020.2964824.

[6] Sean Curtis, Andrew Best, and Dinesh Manocha. Menge: A modular framework for simulating crowd movement. *Collective Dynamics*, 1, 2016. URL https://collective-dynamics.eu/index.php/cod/article/view/A1.

[7] Patrick Dendorfer, Aljoša Ošep, and Laura Leal-Taix'e. Goal-gan: Multimodal trajectory prediction based on goal position estimation. In *Asian Conference on Computer Vision (ACCV)*, 2020.

[8] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. Social GAN: Socially acceptable trajectories with generative adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[9] Peter E. Hart, Nils J. Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968.

[10] Dirk Helbing and Péter Molnár. Social force model for pedestrian dynamics. *Physical Review E*, 51(5): 4282–4286, May 1995. ISSN 1095-3787. URL http://dx.doi.org/10.1103/PhysRevE.51.4282.

[11] Kapil D. Katyal, Gregory D. Hager, and Chien-Ming Huang. Intent-aware pedestrian prediction for adaptive crowd navigation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3277–3283, 2020.

[12] Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton. Imagenet classification with deep convolutional neural networks. *Neural Information Processing Systems*, 25, 01 2012.

[13] Alon Lerner, Yiorgos Chrysanthou, and Dani Lischinski. Crowds by example. *Computer Graphics Forum*, 26(3): 655–664, 2007. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8659.2007.01089.x.

[14] S. Pellegrini, A. Ess, K. Schindler, and L. van Gool. You'll never walk alone: Modeling social behavior for multi-target tracking. In *2009 IEEE 12th International Conference on Computer Vision*, pages 261–268, 2009.

[15] Ashwini Pokle, Roberto Martin-Martin, Patrick Goebel, Vincent Chow, Hans M. Ewald, Junwei Yang, Zhenkai Wang, Amir Sadeghian, Dorsa Sadigh, Silvio Savarese, and et al. Deep local trajectory replanning and control for robot navigation. *2019 IEEE International Conference on Robotics and Automation (ICRA 2019)*. URL http://dx.doi.org/10.1109/ICRA.2019.8794062.

[16] Claudia Pérez-D'Arpino, Can Liu, Patrick Goebel, Roberto Martín-Martín, and Silvio Savarese. Robot navigation in constrained pedestrian environments using reinforcement learning, 2021 IEEE International Conference on Robotics and Automation (ICRA 2021). URL https://arxiv.org/pdf/2010.08600.pdf.

[17] Amir Sadeghian, Vineet Kosaraju, Ali Sadeghian, Noriaki Hirose, Hamid Rezatofighi, and Silvio Savarese. Sophie: An attentive gan for predicting paths compliant to social and physical constraints. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1349–1358, 2019.

[18] Bokui Shen, Fei Xia, Chengshu Li, Roberto Martín-Martín, Linxi Fan, Guanzhi Wang, Claudia Pérez-D'Arpino, Shyamal Buch, Sanjana Srivastava, Lyne P. Tchapmi, Micael E. Tchapmi, Kent Vainio, Li Fei-Fei, and Silvio Savarese. iGibson, a simulation environment for interactive tasks in large realistic scenes, 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2021). URL https://arxiv.org/pdf/2012.02924.pdf.

[19] Lei Tai, Jingwei Zhang, Ming Liu, and Wolfram Burgard. Socially compliant navigation through raw depth inputs with generative adversarial imitation learning. In *2018 IEEE International Conference on Robotics and Automation (ICRA 2018)*.

[20] Jur Van Den Berg, Stephen J Guy, Ming Lin, and Dinesh Manocha. Reciprocal n-body collision avoidance. *Robotics research*, pages 3–19, 2011.