

Occupancy Map Prediction Using Generative and Fully Convolutional Networks for Robot Navigation

Kapil Katyal^{1,2}, Katie Popek¹, Chris Paxton², Joseph Moore¹,
Kevin Wolfe¹, Philippe Burlina¹, and Gregory D. Hager²

Abstract—Fast, collision-free motion through unknown environments remains a challenging problem for robotic systems. In these situations, the robot’s ability to reason about its future motion is often severely limited by sensor field of view (FOV). By contrast, biological systems routinely make decisions by taking into consideration what might exist beyond their FOV based on prior experience. In this paper, we present an approach for predicting occupancy map representations of sensor data for future robot motions using deep neural networks. We evaluate several deep network architectures, including purely generative and adversarial models. Testing on both simulated and real environments we demonstrated performance both qualitatively and quantitatively, with SSIM similarity measure up to 0.899. We showed that it is possible to make predictions about occupied space beyond the physical robot’s FOV from simulated training data. In the future, this method will allow robots to navigate through unknown environments in a faster, safer manner.

I. INTRODUCTION

High-speed, autonomous navigation is predicated on the ability to reason about the environment for effective, collision-free path planning. Existing approaches operate on current sensor readings to update an occupancy map corresponding to an internal representation of where obstacles exist in the environment. These occupancy maps are then used by planning algorithms to generate a collision-free path to a target goal. One of the limitations of this approach is that the planning horizon is limited to the field of view (FOV) of the sensor.

On the other hand, behavioral neuroscience and biological psychology point to the potential role of prediction for navigation in animals and humans. In particular, the hippocampus appears to exhibit some neuronal structures as well as firing sequences that could support not only mapping but also predictive mapping capabilities [2]. Indeed, humans continuously make predictions of what to expect based on past experiences. This allows us to adjust our control policy in real time depending on how close our observations match our predictions [9]. The advantage is most evident while running along a hallway approaching a T-intersection. Even though we cannot see the left or right paths, we generally assume the straight lines will continue and we can predict how the hallway will appear as we turn the corner. Because of this prediction, we do not adjust our running speed unless our prediction is inaccurate. Following this intuition, we believe future predictions of occupancy maps can enable risk sensitive

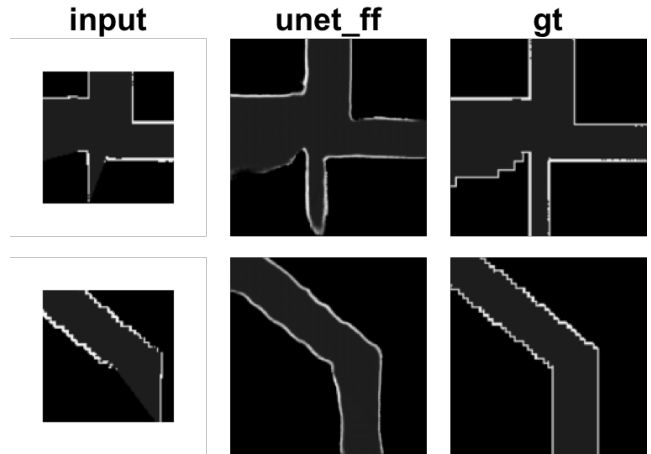


Fig. 1. Two samples of predicted images. The left image is the input based on sensor readings, middle image is the predicted expanded occupancy map (1.50x) and the right image is the ground truth.

control policies for mobile and aerial vehicles. By being able to predict occupancy maps, we can enable faster navigation as the planning horizon can extend beyond the sensor’s limited FOV.

This concept is similar to image completion, a problem for which multiple solutions have been suggested in the past [1, 3]. We take an alternate approach, leveraging the fact that structural information from the observed geometry of the world that can help us make useful predictions about the environment. Deep neural networks have significant advantages over other approaches when used for image completion or image generation [23]. One of the most promising innovations in deep learning has been the development of autoencoder networks and generative adversarial networks (GANs) [5]. These networks use a minimax game adversarial training with opponent generative and discriminative networks, and are capable of encoding a latent representation of images used to generate new examples from latent space.

In this paper, we demonstrate the ability to generate future predictions of occupancy maps without an explicit model using a variety of different neural network architectures with two examples shown in Fig. 1.

The main contributions of this work include:

- A dataset consisting of simulated and physical occupancy maps that can be used to train and validate neural

¹Johns Hopkins University Applied Physics Lab, Laurel, MD, USA. Kapil.Katyal@jhuapl.edu

²Dept. of Comp. Sci., Johns Hopkins University, Baltimore, MD, USA.

- networks for predicting occupancy maps,
- A framework to evaluate the performance and accuracy of different neural network architectures,
- Qualitative and quantitative analysis of the prediction capabilities, performance and accuracy of various network architectures,
- Validation of our approach using occupancy maps generated by a physical LIDAR sensor.

II. RELATED WORK

Model Predictive Control

High-speed navigation has been an active area of research primarily focusing on trajectory optimization, path planning and state estimation. Several papers have investigated model predictive control (MPC) techniques for navigation including [18, 13] however these approaches typically model the vehicle dynamics to predict vehicle motion and not necessarily the environment.

Deep Learning for Generative Models

Deep neural networks have been used in a number of promising ways to achieve high performance in domains such as vision, speech and more recently in robotics manipulation [4, 12].

Oh et. al. used feedforward and recurrent neural networks to perform action-conditional video prediction using Atari games with promising results [16]. These have also been used in image completion, e.g., by Ulyanov et al. [22]. In addition, GANs have demonstrated a promising method for image generation [5]. Isola et al. proposed an approach for training conditional GANs which create one image from another image [8].

Deep Learning for Navigation

More recently, several papers have described approaches to combine elements of deep neural networks with autonomous navigation. These include using deep neural networks for model predictive control [4]. Tamar et al. proposed Value Iteration Networks, which embed a planner inside a deep neural net architecture [21]. Several papers investigate the use of deep reinforcement learning to develop collision-free planning without the need of an internal map, however, these are still restricted by the sensor’s FOV [20, 11].

While each of these papers makes promising contributions to their respective fields, none of the prior works use neural networks and in particular, GANs to generate predictions of future occupancy maps, nor do they focus on extending the planning horizon beyond the sensor’s FOV.

III. PROPOSED ARCHITECTURES

The goal of our network architecture is to learn a function that maps an input occupancy map to an expanded occupancy map that extends beyond the FOV of the sensor. More formally, we are learning the function

$$f : x \rightarrow y_i$$

where x represents the state, in this case, the input occupancy map as an image, y_i represents the output occupancy map and $i \in \mathbb{R}$ represents percent increase of the expanded occupancy map. Components of the function f include an encoding function $f_{enc}(x) \rightarrow h \in \mathcal{H}$ which maps the state space, input occupancy maps to a hidden state and $f_{dec}(h) \rightarrow (y_i)$, which is a decoding function mapping the hidden state to an expanded, predicted occupancy map.

In our experiments, we compare several different neural network architectures including:

- (A) a feedforward network based on a U-Net architecture (**unet_ff**)
- (B) a feedforward network based on the ResNet architecture (**resnet_ff**)
- (C) a GAN using the feedforward network from (a) as the generative network (**gan**)

A. U-Net Feedforward Model

The U-Net feedforward model is based on the network architecture defined by Ronneberger et. al [19] and consists of skip connections which allows a direct connection between layers i and $n - i$ enabling the option to bypass the bottleneck associated with the downsampling layers in order to perform an identity operation. Similar to [8], the encoder network consists of 8 convolution, batch normalization and ReLU layers where each convolution consists of a 4×4 filter and stride length of 2. The number of filters for the 8 layers in the encoder network are: (64, 128, 256, 512, 512, 512, 512, 512). The decoder network consists of 8 upsampling layers with the following number of filters: (512, 1024, 1024, 1024, 1024, 512, 256, 128).

B. ResNet Feedforward Model

The ResNet feedforward model is based on the work by Johnson et. al [10] which consists of 2 convolution layers with stride 2, 9 residual blocks as defined by [6] and two deconvolution layers with with a stride of $\frac{1}{2}$. A key reason this network was selected was based on the ability to learn identify functions, which is key to image translation as well as the success in image-to-image translation demonstrated by the CycleGAN network [24].

C. GAN Model

The GAN networks is based on the pix2pix architecture [8] which has demonstrated impressive results in general purpose image translation including generating street scenes, building facades and aerial images to maps. This network uses the U-Net Feedforward model defined in section III-A and consists of a 6 layer discriminator network with filter sizes: (64, 128, 256, 512, 512, 512).

IV. SIMULATED DATA EXPERIMENTS

Our approach to testing occupancy map prediction using the networks defined above first involved generating a dataset and then performing qualitative and quantitative analysis of the predicted images compared to the ground truth.

A. Data Collection

A dataset of approximately 6000 images of occupancy map subsets was created by simulating a non-holonomic robot moving through a two-dimensional map with a planar LIDAR sensor in C++ with ROS and the OctoMap library [7]. Two maps were created in Solidworks with the path width varying between 3.5 m to 10 m. These were converted into OctoMap’s binary tree format using binvox [14, 15] followed by OctoMap’s binvox2bt tool. The result is an occupancy map with all unoccupied space set as free. We require space outside of the walls to be marked as unknown to provide a ground truth for our estimated maps. These ground truth maps were created by fully exploring the original occupancy maps.

The robot is modeled as a Dubins car, with a state vector $\mathbf{x} = [x, y, \theta]$ and inputs $\mathbf{u} = [v, \dot{\theta}]$ where (x, y) is the robot’s position, v is the velocity, and θ and $\dot{\theta}$ are the heading angle and angular velocity, respectively. For simplicity, the robot is constrained to move at fixed forward velocity of 0.5 m/s. A planar LIDAR sensor with a scanning area of 270° and range of 20 m is used to simulate returns given the robot’s current pose against the ground truth map. These simulated returns are used to create the estimated occupancy map. Path planning is done with nonlinear model-predictive control and direct transcription at 10 Hz. At each time step, a subset of the maps (both the estimated and ground truth) are saved. A 5 m by 5 m square centered around the robot’s pose was chosen with a resolution of 0.05 m. At each time step, the robot’s current state and action space are also logged. Occupancy maps are expanded over time, so our simulation performs a continuous trajectory and the data set is built consecutively instead of randomly sampling throughout a map. A total of six trajectories were simulated. Four paths were used for training data (5221 images) and two were used as a test set (1090 images). Ground truth datasets of the expanded occupancy maps were also generated. These expanded occupancy maps range from 1.10x to 2.00x expansion in increments of 0.10x, e.g., a 2.00x expansion results in a 10 m by 10 m square subset centered around the robot.

B. Training Details

We trained each variant of the neural network using the expanded ground truth occupancy maps from scratch for 200 epochs with a batch size of 1. A total of 15 training sessions were performed to evaluate each of the three neural network architectures across five expansion increases (1.10x, 1.30x, 1.50x, 1.70x, and 2.00x). We use the Adam optimizer with an initial learning rate of 0.0002 and momentum parameters $\beta_1 = 0.5, \beta_2 = 0.999$. In the feedforward models, L1 loss was used as proposed in PatchGan [8] and in the GAN model L1+discriminator loss was used. The decoder layers of the network used a dropout rate of 0.50 and weights were initialized from a Normal distribution ($\mu = 0, \sigma = 0.2$). All models were implemented using PyTorch [17].

C. Simulation Results

Table I provides the structural similarity index metric (SSIM) for each of the networks. Based on the SSIM metric, it can be seen that the U-Net feedforward model outperforms the other networks at 1.10x and 1.30x expansion confirming the qualitative assessment. The quality of the prediction generally decreases as the expansion percentage increases and with expansions 1.50x and above the three networks achieve similar performance.

TABLE I
SSIM ANALYSIS FOR SIMULATION DATA

Expansion	UNET_FF	RESNET_FF	GAN
1.10x	0.899	0.861	0.818
1.30x	0.818	0.780	0.790
1.50x	0.770	0.773	0.759
1.70x	0.760	0.752	0.736
2.00x	0.767	0.770	0.574

V. PHYSICAL EXPERIMENTS

Our next experiment focused on validating our approach with occupancy maps generated by a physical LIDAR sensor.

A. Data Collection and Training Details

In this experiment, we teleoperated a TurtleBot2 robot with a mounted Hokuyo UST-20LX LIDAR sensor around a building. The OctoMap library [7] along with a custom C++ implementation of a particle filter running at 20 Hz was used for simultaneous localization and mapping. The final map was used as ground truth. At each time step a 5 m by 5 m square subset centered around the robot’s current pose of both the ground truth and estimated maps was saved (100 images). Expanded ground truth occupancy maps were generated ranging from 1.10x to 2.00x in 0.10x increments.

Our objective was to evaluate whether training performed on a simulated dataset could be directly transferred to occupancy maps generated by a physical LIDAR sensor. For this reason, we opted to not fine tune the networks using the physical dataset.

B. Physical Results

Fig. 2 represents sample predictions obtained by running the networks trained using simulation data on the occupancy maps generated by the physical sensor. Table II displays the SSIM metric across each of the networks. In the physical experiments, the data is more inconclusive. Similar to the simulation experiments, the quality of the predictions generally decrease as the predicted distance increases, however there was no noticeable difference across the three networks.

VI. CONCLUSION

The ability to perform predictions is key to navigation. This capability is also motivated from the perspective of behavioral neuroscience and psychology. In particular it has been found [2] that certain neuronal structures point to mapping capabilities and may be involved in encoding predictive




Fig. 2. This figure describes the input data, the predicted images and the ground truth for each of the neural networks evaluated on the physical dataset across an expanding prediction window from 1.10x increase to 1.70x increase.

TABLE II
SSIM ANALYSIS FOR PHYSICAL DATA

Expansion	UNET FF	resnet FF	GAN
1.10x	0.511	0.523	0.489
1.30x	0.517	0.534	0.524
1.50x	0.498	0.489	0.511
1.70x	0.486	0.504	0.486
2.00x	0.484	0.499	0.424

mapping events based on past experience. In this paper, our goal is to develop techniques that enable future predictions of occupied space for robotic navigation.

The main intuition behind our predictive approach is that knowledge of the geometry of existing occupied space can serve as a prior for generating predictions. Prior to deep learning, the best methods of generating predictions were through *explicit models*, however, modeling observations and experiences can be difficult if not impossible. Deep learning enables the ability to find hidden representations that encode prior knowledge by collecting datasets that represent experiences. In our work, we leveraged the power of deep learning to encode prior knowledge of likely spatial structures and used this representation to generate future predictions without

an explicit model. We also validated our approach in semi-structured 2D environments and plan to evaluate highly unstructured environments such as caves as future work. While the results presented have primarily focused on 2D navigation, we believe these results can easily extend to 3D by using 3D convolutions.

Our long term objective is to develop risk-sensitive control algorithms capable of leveraging known obstacles in the environment as well as predicted obstacles. In this paper, we have laid the foundation to demonstrate deep networks can be used to make predictions of occupancy maps that extend beyond the FOV of the sensor. In our evaluation, we uncovered conditions where predictions were highly accurate and examples where the predicted results could be improved. We have also verified the inference time of our prediction algorithms average 2.5 ms on a laptop with a NVIDIA GTX 1080 GPU, demonstrating the ability to run these algorithms in real time. As future work, we plan to evaluate prediction mechanisms operating on visual and depth data, to develop an on-line learning policy and also to further develop risk-sensitive control policies for high speed navigation based on these predictions.

REFERENCES

- [1] Marcelo Bertalmio, Guillermo Sapiro, Vincent Caselles, and Coloma Ballester. Image inpainting. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '00*, pages 417–424, New York, NY, USA, 2000. ACM Press/Addison-Wesley Publishing Co. ISBN 1-58113-208-5. doi: 10.1145/344779.344972.
- [2] Randy L Buckner. The role of the hippocampus in prediction and imagination. *Annual review of psychology*, 61:27–48, 2010.
- [3] Tony F. Chan and Jianhong Shen. Mathematical models for local nontexture inpaintings. *SIAM J. Appl. Math.*, 62: 1019–1043, 2002.
- [4] Chelsea Finn and Sergey Levine. Deep visual foresight for planning robot motion. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on*, pages 2786–2793. IEEE, 2017.
- [5] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc., 2014. URL <http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>.
- [6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015. URL <http://arxiv.org/abs/1512.03385>.
- [7] Armin Hornung, Kai M. Wurm, Maren Bennewitz, Cyrill Stachniss, and Wolfram Burgard. OctoMap: An efficient probabilistic 3D mapping framework based on octrees. *Autonomous Robots*, 2013. doi: 10.1007/s10514-012-9321-0. URL <http://octomap.github.com>.
- [8] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. *CVPR*, 2017.
- [9] Michael Joch, Mathias Hegele, Heiko Maurer, Hermann Miller, and Lisa Katharina Maurer. Brain negativity as an indicator of predictive error processing: the contribution of visual action effect monitoring. *Journal of Neurophysiology*, 118(1):486–495, 2017. doi: 10.1152/jn.00036.2017. URL <https://doi.org/10.1152/jn.00036.2017>. PMID: 28446578.
- [10] Justin Johnson, Alexandre Alahi, and Fei-Fei Li. Perceptual losses for real-time style transfer and super-resolution. *CoRR*, abs/1603.08155, 2016. URL <http://arxiv.org/abs/1603.08155>.
- [11] Gregory Kahn, Adam Villafior, Bosen Ding, Pieter Abbeel, and Sergey Levine. Self-supervised deep reinforcement learning with generalized computation graphs for robot navigation. *CoRR*, abs/1709.10489, 2017. URL <http://arxiv.org/abs/1709.10489>.
- [12] Sergey Levine, Peter Pastor, Alex Krizhevsky, and Deirdre Quillen. Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection. *arXiv preprint arXiv:1603.02199*, 2016.
- [13] S. Maniatopoulos, D. Panagou, and K. J. Kyriakopoulos. Model predictive control for the navigation of a nonholonomic vehicle with field-of-view constraints. In *2013 American Control Conference*, pages 3967–3972, June 2013. doi: 10.1109/ACC.2013.6580446.
- [14] Patrick Min. <http://www.patrickmin.com/binvox>, 2004 - 2017. Accessed: 2017-02-20.
- [15] Fakir S. Nooruddin and Greg Turk. Simplification and repair of polygonal models using volumetric techniques. *IEEE Transactions on Visualization and Computer Graphics*, 9(2):191–205, 2003.
- [16] Junhyuk Oh, Xiaoxiao Guo, Honglak Lee, Richard L Lewis, and Satinder Singh. Action-Conditional Video Prediction using Deep Networks in Atari Games. In C Cortes, N D Lawrence, D D Lee, M Sugiyama, R Garnett, and R Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 2845–2853. Curran Associates, Inc., 2015.
- [17] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- [18] J. L. Piovesan and H. G. Tanner. Randomized model predictive control for robot navigation. In *2009 IEEE International Conference on Robotics and Automation*, pages 94–99, May 2009. doi: 10.1109/ROBOT.2009.5152468.
- [19] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. *CoRR*, abs/1505.04597, 2015. URL <http://arxiv.org/abs/1505.04597>.
- [20] Lei Tai, Giuseppe Paolo, and Ming Liu. Virtual-to-real deep reinforcement learning: Continuous control of mobile robots for mapless navigation. *CoRR*, abs/1703.00420, 2017. URL <http://arxiv.org/abs/1703.00420>.
- [21] Aviv Tamar, Yi Wu, Garrett Thomas, Sergey Levine, and Pieter Abbeel. Value iteration networks. In *Advances in Neural Information Processing Systems*, pages 2154–2162, 2016.
- [22] Dmitry Ulyanov, Andrea Vedaldi, and Victor S. Lempitky. Deep image prior. *CoRR*, abs/1711.10925, 2017. URL <http://arxiv.org/abs/1711.10925>.
- [23] Raymond A. Yeh, Chen Chen, Teck-Yian Lim, Mark Hasegawa-Johnson, and Minh N. Do. Semantic image inpainting with perceptual and contextual losses. *CoRR*, abs/1607.07539, 2016. URL <http://arxiv.org/abs/1607.07539>.
- [24] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. *CoRR*, abs/1703.10593, 2017. URL <http://arxiv.org/abs/1703.10593>.