Benchmarking Augmentation Methods for Learning Robust Navigation Agents: The Winning Entry of the 2021 iGibson Challenge

Naoki Yokoyama[1*] Qian Luo[1*] Dhruv Batra[1,2] Sehoon Ha[1]

[1]Georgia Tech [2]Meta AI

{nyokoyama,luoqian,dbatra,sehoonha}@gatech.edu

Abstract

While impressive progress has been made for teaching embodied agents to navigate static environments using vision, much less progress has been made on more dynamic environments that may include moving pedestrians or movable obstacles. In this study, we aim to benchmark different augmentation techniques for improving the agent's performance in these challenging environments. We show that adding several dynamic obstacles into the scene during training confers significant improvements in test-time generalization, achieving much higher success rates than baseline agents. We find that this approach can also be combined with image augmentation methods to achieve even higher success rates. Additionally, we show that this approach is also more robust to sim-to-sim transfer than image augmentation methods. Finally, we demonstrate the effectiveness of this dynamic obstacle augmentation approach by using it to train an agent for the 2021 iGibson Challenge at CVPR, achieving 1st place for Interactive Navigation. An extended version of this work can be found at http://arxiv.org/abs/2109.10493.

1. Introduction

Mobile robots must be able to skillfully navigate through their environments to operate effectively in the real world. Fortunately, several recent works using deep reinforcement learning have shown promising results by deploying robots that can successfully navigate in novel environments in the real world [3, 7, 8]. However, datasets and simulators for navigation tasks featuring dynamic objects are not as abundant as the static environments used by these works.

To address the challenges of dynamic navigation and constraints on the availability of training environments, we aim to provide a benchmark and analysis of augmentation techniques for visual navigation such that other researchers can leverage our findings from the iGibson Challenge. In particular, we show that simply adding several dynamic obstacles (pedestrians) to the scene during training improves the agent's test-time success rate in novel environments significantly, even for navigation tasks that *do not involve dynamic obstacles in the environment*, such as PointNav and InteractiveNav. We conduct a systematic analysis in which we sweep through both the amount of training data available and the amount of pedestrians used during training.

Figure 1. Dynamic obstacle augmentation adds several moving pedestrians as obstacles (visualized in RGB, left, and depth, right).

We then compare this dynamic obstacle augmentation method against two image augmentation methods, *Crop* and *Cutout*, which was shown by Laskin et al. [4] to significantly improve test-time generalization (up to 4x performance improvement) on various benchmarks for visual tasks [2, 5]. We show that dynamic obstacle augmentation can be combined synergistically with image augmentations to further improve performance. Similar to Laskin et al., we also show that in contrast, combining different image augmentation methods can reduce performance gains.

We demonstrate the effectiveness of dynamic obstacle augmentation by participating in the 2021 iGibson Challenge at CVPR [6], where our agent ranked 1st place for InteractiveNav. This feat was accomplished without using any rewards specific to InteractiveNav; only a basic PointNav reward and dynamic obstacle augmentation was used. Despite this, our approach achieved a 4% (absolute percentage) higher success rate than the 2nd place team.

*Indicates equal contribution.

2. Augmentation Techniques

To improve the test-time generalization of the agent for visual navigation, we introduce several moving pedestrians into the environment (illustrated in Figure 1). This approach aims to prevent the agent from learning to overfit to the environments it sees during training; with a large number of moving visual distractors present in the scene, it is more difficult for the agent to memorize the layout of the training environment using observations from its camera. Additionally, unlike image augmentation methods that simply perturb the agent's visual inputs, this method forces the agent to learn a larger variety of paths for a given episode's pair of start-goal positions. Because the paths that the pedestrians take are newly generated each time the environment is reset, the agent is not able to always take the same path without colliding into a pedestrian.

We compare this method against two image augmentation methods that have been shown to improve performance for reinforcement learning using visual data [4]. The *Crop* image augmentation extracts a random patch from the original frame; similar to Laskin et al. [4], we crop the frame at a random location such that the resulting image is 8% shorter in height and width. *Cutout*, detailed by Zhong et al. [9], inserts a black rectangle of a random shape, aspect ratio, and location into the original frame. In addition, we investigate *Crop&Cutout*, which combines *Crop* and *Cutout* by applying them sequentially.

3. Results

Table 1. **PointNav (0 ppl)** Success Rates

# of train scenes	# of dynamic pedestrians during training (baseline) 0 ppl	3 ppl	6 ppl	12 ppl	18 ppl
1	$74.93_{\pm 1.91}$	$73.90_{\pm 3.26}$	$\mathbf{78.64}_{\pm 1.76}$	$76.53_{\pm 2.65}$	$78.12_{\pm 0.57}$
2	$81.41_{\pm 2.03}$	$80.85_{\pm 1.30}$	$81.22_{\pm 0.89}$	$81.69_{\pm 1.17}$	$\mathbf{82.86}_{\pm 1.26}$
4	$82.39_{\pm 0.87}$	$\mathbf{86.15}_{\pm 0.57}$	$84.37_{\pm 1.79}$	$84.32_{\pm 2.07}$	$85.77_{\pm 0.61}$
8	$87.89_{\pm 1.31}$	$\mathbf{91.41}_{\pm 0.80}$	$88.97_{\pm 1.65}$	$90.99_{\pm 1.52}$	$89.44_{\pm 0.34}$
16	$91.60_{\pm 1.68}$	$90.99_{\pm 0.98}$	$91.69_{\pm 0.53}$	$\mathbf{92.07}_{\pm 0.81}$	$91.69_{\pm 0.70}$
32	$90.80_{\pm 0.78}$	$92.68_{\pm 0.40}$	$\mathbf{93.90}_{\pm 0.65}$	$92.68_{\pm 1.33}$	$92.39_{\pm 0.69}$
72	$94.84_{\pm 1.50}$	$\mathbf{95.26}_{\pm 0.46}$	$93.43_{\pm 1.16}$	$94.93_{\pm 0.83}$	$94.08_{\pm 0.61}$

We evaluate performance while varying both the amount of available training scenes and the number of dynamic pedestrians used during training. For these experiments, we tackle two different tasks: PointNav and SocialNav. PointNav, as detailed in [1], tasks the robot to navigate from a start position to a goal coordinate in the environment. SocialNav is a variant of PointNav, in which there are moving pedestrians in the environment that the robot cannot collide with. We did not evaluate agents trained with no pedestrians for SocialNav due to their inability to adapt to dynamic pedestrians at test-time, which led to poor performance.

Table 2. **SocialNav (3 ppl)** Success Rates

# of train scenes	# of dynamic pedestrians during training (baseline) 3 ppl	6 ppl	12 ppl	18 ppl
1	$65.26_{\pm 2.47}$	$\mathbf{68.78}_{\pm 0.35}$	$65.40_{\pm 3.77}$	$68.40_{\pm 2.26}$
2	$68.97_{\pm 0.96}$	$69.20_{\pm 2.92}$	$\mathbf{72.68}_{\pm 2.03}$	$71.31_{\pm 1.66}$
4	$74.60_{\pm 1.45}$	$\mathbf{77.32}_{\pm 0.64}$	$76.10_{\pm 0.24}$	$75.77_{\pm 1.52}$
8	$80.38_{\pm 1.27}$	$\mathbf{80.89}_{\pm 1.04}$	$79.91_{\pm 3.37}$	$78.64_{\pm 1.04}$
16	$83.00_{\pm 1.37}$	$\mathbf{83.71}_{\pm 0.59}$	$82.96_{\pm 1.61}$	$80.47_{\pm 0.63}$
32	$84.04_{\pm 1.09}$	$\mathbf{84.41}_{\pm 0.74}$	$83.99_{\pm 1.29}$	$81.27_{\pm 0.87}$
72	$85.26_{\pm 1.27}$	$84.79_{\pm 1.33}$	$\mathbf{86.43}_{\pm 1.75}$	$83.10_{\pm 2.93}$

PointNav Success

	Dyn.	Crop	Cutout
Dyn.	88.97	-	-
Crop	91.5	87.32	-
Cutout	91.69	90.28	91.55

No augmentation (0 ppl): 87.89%

SocialNav Success

	Dyn.	Crop	Cutout
Dyn.	80.89	-	-
Crop	82.82	81.78	-
Cutout	82.3	79.93	82.44

No augmentation (3 ppl): 80.83%

Figure 2. Heatmaps comparing average success rates over three random seeds when trained with various augmentation methods. Adding *Dyn* leads *Crop* to attain the highest success rates for both tasks and does not decrease the performance of *Cutout*.

Our results indicate that dynamic obstacle augmentation improves performance for both visual navigation tasks. Both tables show performance gains provided by dynamic obstacle augmentation are more substantial when only a small number (8 or less) of training scenes are available. As the number of available training scenes increases, the gap between the baseline agents and those trained with added dynamic pedestrians decreases. This indicates that the limited number of training scenes causes overfitting, but can be mitigated by dynamic obstacle augmentation.

We compare our technique against the *Crop* and *Cutout* image augmentation techniques. To narrow the scope of our investigation, we limit these experiments to agents trained with only eight scenes. When dynamic obstacle augmentation is used, six pedestrians are added during training. We find that using *Cutout* alone outperforms using dynamic obstacle augmentation alone by a slight margin for both tasks. We also evaluate the effects of combining the image augmentation techniques. We observe that combining *Cutout* and *Crop* (bottom-middle cells) only leads to worse performance than simply using *Cutout* alone. However, as shown in Figure 2, when dynamic obstacle augmentation is combined with either the *Crop* or *Cutout* image augmentations, we find that the success rate either remains the same or is further improved. Combining *Dyn* with *Crop* (middle-left cells) yields a higher success rate than using either method alone, and achieves the highest success rate for both tasks, tying with *Cutout* for PointNav and surpassing it for SocialNav. Combining *Cutout* with *Dyn* (bottom-left cells) does not affect performance over using *Cutout* alone.

References

[1] Peter Anderson, Angel Chang, Devendra Singh Chaplot, Alexey Dosovitskiy, Saurabh Gupta, Vladlen Koltun, Jana Kosecka, Jitendra Malik, Roozbeh Mottaghi, Manolis Savva, et al. On evaluation of embodied navigation agents. *arXiv preprint arXiv:1807.06757*, 2018. 2

[2] Karl Cobbe, Chris Hesse, Jacob Hilton, and John Schulman. Leveraging procedural generation to benchmark reinforcement learning. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 2048–2056. PMLR, 13–18 Jul 2020. 1

[3] Abhishek Kadian, Joanne Truong, Aaron Gokaslan, Alexander Clegg, Erik Wijmans, Stefan Lee, Manolis Savva, Sonia Chernova, and Dhruv Batra. Sim2real predictivity: Does evaluation in simulation predict real-world performance? *IEEE Robotics and Automation Letters (RA-L)*, 2020. 1

[4] Michael Laskin, Kimin Lee, Adam Stooke, Lerrel Pinto, Pieter Abbeel, and Aravind Srinivas. Reinforcement learning with augmented data. *arXiv preprint arXiv:2004.14990*, 2020. 1, 2

[5] Michael Laskin, Aravind Srinivas, and Pieter Abbeel. Curl: Contrastive unsupervised representations for reinforcement learning. *Proceedings of the 37th International Conference on Machine Learning, Vienna, Austria, PMLR 119*, 2020. arXiv:2004.04136. 1

[6] Chengshu Li, Jaewoo Jang, Fei Xia, Roberto Martín-Martín, Claudia D'Arpino, Alexander Toshev, Anthony Francis, Edward Lee, and Silvio Savarese. iGibson Challenge 2021, 2021. 1

[7] Ruslan Partsey, Erik Wijmans, Naoki Yokoyama, Oles Dobosevych, Dhruv Batra, and Oleksandr Maksymets. Is mapping necessary for realistic pointgoal navigation? In *Computer Vision and Pattern Recognition (CVPR), 2022 IEEE Conference on*. IEEE, 2022. 1

[8] Naoki Yokoyama, Sehoon Ha, and Dhruv Batra. Success weighted by completion time: A dynamics-aware evaluation criteria for embodied navigation. 2021. 1

[9] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. Random erasing data augmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13001–13008, 2020. 2