Searching for Problematic Simulation Conditions

Elizabeth Johnson
Computer Science Department
Pomona College
Claremont, California, USA
emja2020@mymail.pomona.edu

Simon Heck
Computer Science Department
Pomona College
Claremont, California, USA
simonpheck@gmail.com

Keneth Gonzalez Hernandez
Computer Science Department
Pomona College
Claremont, California, USA
kmga2020@mymail.pomona.edu

Anthony J. Clark
Computer Science Department
Pomona College
Claremont, California, USA
anthony.clark@pomona.edu

***Abstract*—Many robot use-cases put a robot in close contact with people. These scenarios require the robot to make complex decisions that—above all else—must be safe. Often, sensor processing and decision making methods rely heavily on machine learning, and these techniques are only as useful as the training dataset. Current methods and datasets do not account for enough variation or "extraordinary" conditions. We propose using novelty search to discover scenarios causing a model to behave poorly.**

***Index Terms*—robotics, simulation, learning, vision**

Robots are poised to enter use in a wide-range of unstructured and human-centric environments. Cleaning, structural inspection, mining, and giving tours are a few commonly discussed tasks. Take hospital room maintenance, where a robot could aid in the disposal of single-use items and disinfecting surfaces, as an example. UVD Robotics created autonomous mobile robots with UV disinfecting technology to clean patient rooms [1]. Use of the robots in Croatia and Italy showed less hospital staff testing positive for COVID-19 and reducing manual labor spent on disinfecting [2]. Robot performance was also better than that of humans: micro-resistant organisms were present $10\,\%$ of the time in manual disinfecting versus no organisms detected in robot disinfecting. Autonomous disinfecting robots could also be deployed in high-contact public areas such as public transit, bathrooms, and schools.

Robots needing complex sensor processing often rely on machine learning methods. To build such machines, engineers often rely on simulation, which is faster, safer, and more cost effective. A common approach involves: **(1)** developing a simulation environment, **(2)** generating a large dataset, **(3)** training an artificial intelligence model, and then **(4)** tweaking the model so that it works on a real system. The drawback of simulation is that we cannot perfectly model all aspects of reality, and we are always left with a simulation-designed system that may fail when put into a real environment [3].

To partially address the reality-gap, we propose the use of photogrammetry [4], neural radiance fields (NeRF) [5], [6], style transfer [7], and Unreal Engine 5 (a recently released system for creating real-time 3D content and experiences) [8] to create high-fidelity virtual environments. The top row of Figure 1 shows an interior scene modeled using these techniques. These technologies have only recently become usable for roboticists, and open new avenues for generating data and validating results. One can build realistic environments and incorporate virtual creatures with dynamic behaviors.

Existing work focuses on collecting *realistic* training data [9] and introducing random but realistic noise [10]. Such models often perform well in controlled environments but poorly when confronting novel scenarios. How will a trained model react when a child unexpectedly jumps in front of a robot from around a corner—a scenario that could be overlooked during data generation. The drawbacks of realistic data with noise are like those in which systems are trained primarily on datasets comprising images with mostly white people [11] or mostly western imagery [12]: *the datasets perform extremely well on the training datasets, but they fail when tested in real-world scenarios*.

Differing from existing work, we focus on collecting training data from a diversity of scenarios. We create robot training data by creating virtual environments using common techniques and adding: **(1)** people of varying skin colors, styles, appearances (including unrealistically tall or proportioned), and behaviors; **(2)** unexpected changes to surfaces (e.g., changing an interior floor from tile into grass or a ceiling from paint into gravel); **(3)** weather effects (e.g., viewing rain through a window, and window blowing through trees); **(4)** local and exotic flora and fauna; **(5)** erratic lighting (e.g., instantly changing the time or day and day of year); and **(6)** structural deformations (e.g., distorting the shape of a hallway). Examples of randomized lighting and texture are depicted in Figure 1. We will use novelty search [13] to discover scenarios that cause a robot to fail (e.g., by running into a person) and then feed these scenarios back into the data generation process.

The proposed techniques should improve safety by eliminating some unexpected robot behaviors, but they must also be combined with techniques from control theory [14] and software engineering [15] for higher confidence.

Increased automation is likely to eliminate low-wage jobs and impact those already most affected by fluctuations in the economy [16]. So, while eliminating dangerous and monotonous jobs may be *good* in the long term, we must have a plan for those being displaced.

Acknowledgment

The authors gratefully thank Pomona College for supporting this work.

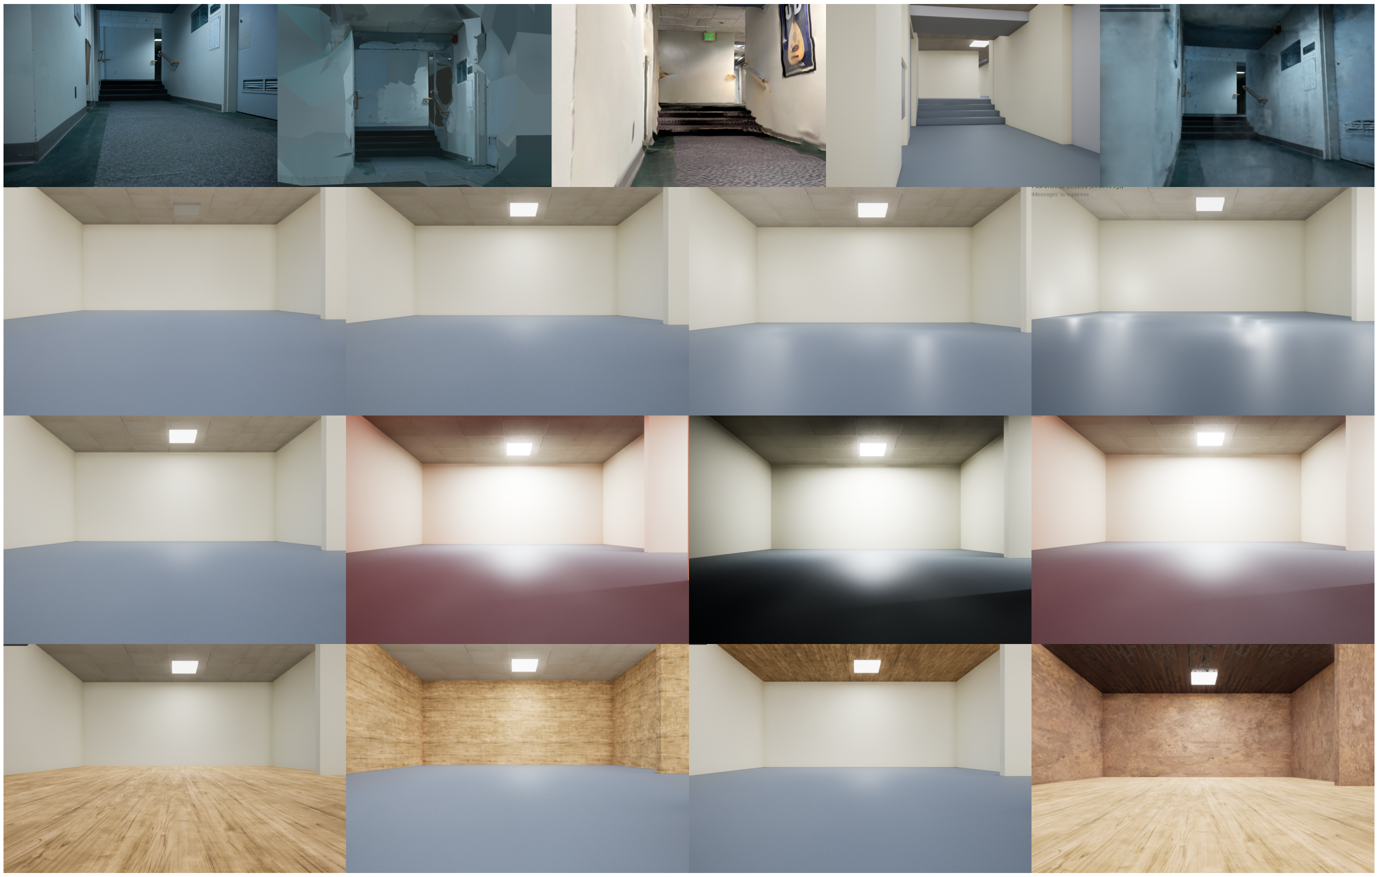

Fig. 1: The first row depicts (from left to right) a photograph, photogrammetry model created with RealityCapture, photogrammetry model created with PolyCam, manually constructed Unreal Engine level, and NeRF model of the same environment. Subsequent rows depict a different view of the same scene, but with randomized lighting and textures.

REFERENCES

[1] E. Ackerman, "Autonomous robots are helping kill coronavirus in hospitals," *IEEE Spectrum*, 2020.

[2] "European commission procurement triggers wider deployment of innovative disinfection robots across europe's hospitals."

[3] J. Truong, S. Chernova, and D. Batra, "Bi-directional Domain Adaptation for Sim2Real Transfer of Embodied Navigation Agents," *arXiv:2011.12421 [cs]*, Nov. 2020.

[4] J. L. Schönberger and J.-M. Frahm, "Structure-from-Motion Revisited," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2016, pp. 4104–4113.

[5] T. Müller, A. Evans, C. Schied, and A. Keller, "Instant neural graphics primitives with a multiresolution hash encoding," *ACM Transactions on Graphics*, vol. 41, no. 4, pp. 1–15, Jul. 2022.

[6] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "NeRF: Representing scenes as neural radiance fields for view synthesis," *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, Jan. 2022.

[7] F. Luan, S. Paris, E. Shechtman, and K. Bala, "Deep Photo Style Transfer," in *2017 IEEE Conference on Computer Vision and Pattern Recognition*. Honolulu, HI: IEEE, Jul. 2017, pp. 6997–7005.

[8] W. Qiu, F. Zhong, Y. Zhang, S. Qiao, Z. Xiao, T. S. Kim, and Y. Wang, "UnrealCV: Virtual Worlds for Computer Vision," in *Proceedings of the 25th ACM International Conference on Multimedia*. Mountain View California USA: ACM, Oct. 2017, pp. 1221–1224.

[9] W. Guerra, E. Tal, V. Murali, G. Ryou, and S. Karaman, "FlightGoggles: Photorealistic Sensor Simulation for Perception-driven Robotics using Photogrammetry and Virtual Reality," *arXiv:1905.11377 [cs]*, May 2019.

[10] S. Zada, I. Benou, and M. Irani, "Pure Noise to the Rescue of Insufficient Data: Improving Imbalanced Classification by Training on Random Noise Images," Jun. 2022.

[11] S. Shankar, Y. Halpern, E. Breck, J. Atwood, J. Wilson, and D. Sculley, "No Classification without Representation: Assessing Geodiversity Issues in Open Data Sets for the Developing World," Nov. 2017.

[12] K. Yang, K. Qinami, L. Fei-Fei, J. Deng, and O. Russakovsky, "Towards fairer datasets: Filtering and balancing the distribution of the people subtree in the ImageNet hierarchy," in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, ser. FAT* '20. New York, NY, USA: Association for Computing Machinery, Jan. 2020, pp. 547–558.

[13] K. O. Stanley and J. Lehman, *Why Greatness Cannot Be Planned: The Myth of the Objective*. Cham, Switzerland: Springer International Publishing, 2015.

[14] H. Tsukamoto, S.-J. Chung, and J.-J. E. Slotine, "Contraction theory for nonlinear stability analysis and learning-based control: A tutorial overview," *Annual Reviews in Control*, vol. 52, pp. 135–169, 2021. [Online]. Available: http://arxiv.org/abs/2110.00675

[15] M. A. Langford, K. H. Chan, J. E. Fleck, P. K. McKinley, and B. H. Cheng, "MoDALAS: Model-driven assurance for learning-enabled autonomous systems," in *2021 ACM/IEEE 24th International Conference on Model Driven Engineering Languages and Systems (MODELS)*, 2021, pp. 182–193.

[16] W. Dauth, S. Findeisen, J. Suedekum, and N. Woessner, "The adjustment of labor markets to robots," *Journal of the European Economic Association*, vol. 19, no. 6, pp. 3104–3153, 2021. [Online]. Available: https://doi.org/10.1093/jeea/jvab012