# The Pervasiveness of Deep Learning in Robotics Research Does Not Impede Scientific Insights into Robotics Problems

Stewart Jamieson*

## I. Introduction

In 2018, papers on research into "Deep Learning for robotics and automation" overtook every other category of papers presented at ICRA [1], despite any apportionment by authors of such papers to burgeoning conferences like CoRL. However, deep learning is not a traditional robotics topic; it extends earlier neural network approaches developed by computer scientists to solve more difficult problems, and only recently became widespread in 2012 when it started a "revolution" in computer vision [2]. While it is natural that roboticists would adapt the leading techniques in computer vision for robot vision applications, in the last several years deep learning methods have been studied extensively for applications in nearly every major sub-domain of robotics, such as localization [3], action planning [4], grasping and manipulation [5], and optimal control for locomotion [6].

The popularity of deep learning among roboticists is not evidence that it is a fruitful path to a greater scientific understanding into open problems in robotics. Nonetheless, I will argue that the pervasiveness of deep learning in robotics research is not an impediment to scientific insights into robotics problems, but rather leads to a probabilistic and data-driven perspective that will be essential for solving the next generation of robotics challenges.

## II. Rebuttal of Counter Arguments

This section disputes claims that deep learning methods give inferior, if any, insights into robotics problems, when compared to traditional approaches.

### A. Deep Learning Systems Can Be Interpretable and Will Lead to New Scientific Insights into Robotics

One argument against encouraging roboticists to conduct deep learning research states that no scientific insights can be made into a robotics problem by applying a deep learning method to solve it. This argument is driven by the belief that deep learning methods are uninterpretable; that their intermediary operations transform data through internal representations that cannot provide scientific insights into the problem or its solution. On the contrary, even an end-to-end deep learning approach can produce meaningful internal structures and representations, such as the road edge detector learned by a deep learning model for autonomous driving using only unprocessed video inputs [7]. This should make techniques for interpreting deep networks of interest to roboticists; deep features may give insights into problems which decades of development of model-based algorithms have not. Deep neural network interpretability is a highly topical subject, with papers from a recent AAAI workshop[1] on it giving insights such as: why deep networks pay more attention to specific parts images [8] and the purposes of individual neurons in vision systems [9]. These insights teach us almost as much about vision tasks as about deep learning.

The study of interpreting deep models and features is mostly absent from robotics literature. While the features learned by deep models for image classification and NLP tasks are extensively studied by computer scientists, it is difficult to find publications examining features learned by models trained for robotics problems like manipulation, control, and localization. This may be due to the difficulty in comprehending the features and feature spaces learned by techniques such as reinforcement learning, but there has been recent work into visualizing and understanding some such features and learning processes [10]. The outputs of deep reinforcement learning, policies, seem to be more easily interpreted; roboticists may soon learn from policies learned in robotics applications just as professional Go players are learning from the strategies of DeepMind's AlphaGo [11].

### B. Deep Learning Research Does Not Impede Research into Developing Safe and Versatile Robots

Another criticism of deep learning among roboticists has been that such systems lack provable robustness guarantees and versatility. Some may argue that deep learning research produces robots that perform well in certain environments but have unknown failure modes and are limited to specific tasks; if these issues were insurmountable, time spent researching deep methods might be better spent improving the performance of more versatile and well understood models. However, one should not assume these issues cannot be solved through studying different model structures or learning strategies. Some successes in producing calibrated uncertainty estimates for deep neural network outputs [12], and developing deep models that can be easily transferred to new tasks and robots [13], suggest it may eventually be possible to develop deep learning systems with probabilistic safety guarantees and high versatility. Additionally, this criticism neglects considering other methods of incorporating deep learning into robots; for example, neural networks have been used to learn optimal control policies while guaranteeing safety and stability [14].

*S. Jamieson is in the MIT-WHOI Joint Ph.D. Program in Applied Ocean Science and Engineering `sjamieson@{mit,whoi}.edu`

[1]AAAI-19 Workshop on Network Interpretability for Deep Learning

## III. Arguments Supporting the Role of Deep Learning in Robotics Research

### A. Data-driven Behaviors Are as Essential as Model-driven Behaviors for Developing Intelligent Robots

There is an old rivalry between data-driven (e.g. deep learning) and model-driven (rule based) approaches in the field of AI; however, in a robotic system both data-driven and model-driven behaviours are essential. One goal of roboticists is to construct robots which obey specific rules but also learn new behaviours online. This saves roboticists from needing to model every new task, and leads to versatile robots that autonomously learn behaviours difficult to describe through rules (e.g. picking up an object with unique geometry). A data-driven approach is ideal for developing robots that learn through demonstration, observation, or trial and error. Research has also shown that data-driven approaches can create systems which better agree with human intuition and behaviour than handcrafted models [15], suggesting they could improve human-robot interaction.

Many new scientific insights are required to fully leverage deep learning in robotic systems. For example, what features of a new task or robot affect the difficulty of re-purposing a pretrained deep model for it? Answering such questions is a role suited to roboticists deeply familiar with the tasks, physics, and systems involved, motivating the involvement of roboticists in deep learning research. Furthermore, the answers to many such questions will be easier to identify after observing how well different deep learning approaches solve a variety of robotics problems, and these observations are being made daily thanks to the pervasiveness of deep learning research among roboticists.

### B. Deep Learning Gives Insights into the Complex Probabilistic Nature of Robot Interactions with the World

A major advancement in robotics was the adoption of probabilistic modelling techniques such as the particle filter [16]. While roboticists have physical models for many processes, such as locomotion, the uncertainty involved in the actual execution of most processes is too complicated to model analytically. Probabilistic modelling techniques better approximate a robot's state given a simple process model of how it changes based on time and actions. These probabilistic models led to many robotics insights into areas like active learning, in which a robot takes actions specifically to reduce its own uncertainty about its own state or the world so that it can operate more robustly.

Deep learning techniques solve similar problems to these probabilistic models in far more complex settings. For example, supervised deep learning systems replace the need for a model of a process with the need for labelled input/output pairs, and learn a good approximation of the process; this is useful when a good model cannot easily be described analytically. Some work has demonstrated that these systems can estimate uncertainties inherent to the modelled process [12], and they can also be used to estimate the sensitivity of the process outputs to input variations. Better understanding the probabilistic nature of robotic processes will provide insights into robust vision and actuation systems, among others. One application of this was an autonomous vehicle that could estimate its own uncertainty in highly complex tasks, such as predicting another vehicles motion, and use this uncertainty to inform both its decisions and the passengers [17].

## IV. Conclusions

It is argued that deep learning research can provide scientific insights into robotics problems, and is in fact essential to enable and study new data-driven behaviours in robots and to understand the uncertainty in more complex interactions between a robot and its environment. Much of this research, such as understanding the learned features and uncertainty estimates, requires specific domain knowledge regarding the tasks, robots, and physics involved, so these scientific insights are best gained through encouraging roboticists to explore and pursue deep learning research.

## References


[1] GVU Center at Georgia Tech, "Workbook: ICRA 2018," 2018.

[2] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 12 2015.

[3] V. Peretroukhin and J. Kelly, "DPC-Net: Deep Pose Correction for Visual Localization," *IEEE Robotics and Automation Letters*, 2018.

[4] P. Karkus, D. Hsu, and W. S. Lee, "QMDP-Net: Deep Learning for Planning under Partial Observability," in *Neural Information Processing Systems*, 2017.

[5] S. Levine, P. Pastor, A. Krizhevsky, J. Ibarz, and D. Quillen, "Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection," *The International Journal of Robotics Research*, vol. 37, no. 4-5, pp. 421–436, 2018.

[6] Y. Chen, Y. Shi, and B. Zhang, "Optimal Control Via Neural Networks: A Convex Approach," in *ICLR 2019*, 2019, pp. 1–25.

[7] M. Jaritz, R. de Charette, M. Toromanoff, E. Perot, and F. Nashashibi, "End-to-End Race Driving with Deep Reinforcement Learning," *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2070–2075, 2018.

[8] S. Xie, D. Chen, R. Zhang, and H. Xue, "Deep Features Analysis with Attention Networks," in *AAAI-19 Workshop on Network Interpretability for Deep Learning*, 2019.

[9] D. Bau, J.-Y. Zhu, H. Strobelt, B. Zhou, J. B. Tenenbaum, W. T. Freeman, and A. Torralba, "Visualizing and Understanding Generative Adversarial Networks (Extended Abstract)," in *AAAI 2019 Workshop on Network Interpretability for Deep Learning*, 2019.

[10] J. Wang, L. Gou, H. W. Shen, and H. Yang, "DQNViz: A Visual Analytics Approach to Understand Deep Q-Networks," *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 1, pp. 288–298, 2019.

[11] V. Lin, "A new perspective on the star point and the implications thereof," 2017.

[12] A. Kendall and Y. Gal, "What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision?" in *Advances in Neural Information Processing Systems 30*, 3 2017, pp. 5574–5584.

[13] C. Devin, A. Gupta, T. Darrell, P. Abbeel, and S. Levine, "Learning modular neural network policies for multi-task and multi-robot transfer," in *Proceedings - IEEE International Conference on Robotics and Automation*, 2017, pp. 2169–2176.

[14] F. Berkenkamp, M. Turchetta, A. P. Schoellig, and A. Krause, "Safe Model-based Reinforcement Learning with Stability Guarantees," in *Neural Information Processing Systems*, 2017.

[15] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The Unreasonable Effectiveness of Deep Features as a Perceptual Metric," in *Computer Vision and Pattern Recognition*, 2018.

[16] S. Thrun, "Particle Filters in Robotics," in *Proceedings of Uncertainty in AI (UAI)*, 2002, pp. 511–518.

[17] R. McAllister, Y. Gal, A. Kendall, M. V. D. Wilk, A. Shah, R. Cipolla, and A. Weller, "Concrete Problems for Autonomous Vehicle Safety: Advantages of Bayesian Deep Learning," in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2017.