

Efficient Adaptation for End-to-End Vision-Based Robotic Manipulation

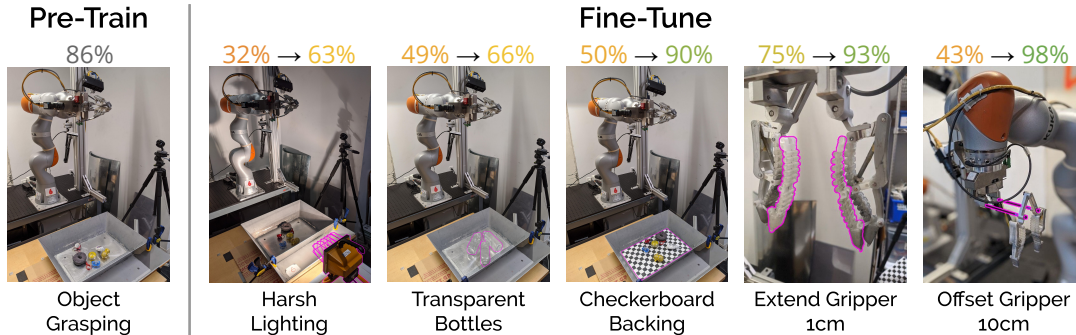


Fig. 1: Original robot configuration used for pre-training (left), and adaptation challenges (highlighted in pink) studied in this work (right) with associated performance improvements (top) obtained using our fine-tuning method.

Abstract—One of the great promises of robot learning systems is that they will be able to learn from their mistakes and continuously adapt to ever-changing environments. Despite this potential, most of the robot learning systems today are deployed as a fixed policy and they are not being adapted after their deployment. In this paper, we present a method and empirical evidence towards a robot learning framework that facilitates continuous adaption. In particular, we demonstrate how to adapt vision-based robotic manipulation policies to new variations by fine-tuning via off-policy reinforcement learning, including changes in background, object shape and appearance, lighting conditions, and robot morphology. Further, this adaptation uses less than 0.2% of the data necessary to learn the task from scratch. We find that the simple approach of fine-tuning pre-trained policies leads to substantial performance gains over the course of fine-tuning, and that pre-training via RL is essential: training from scratch or adapting from supervised ImageNet features are both unsuccessful with such small amounts of data. We also find that these positive results hold in a limited continual learning setting, in which we repeatedly fine-tune a single lineage of policies using data from a succession of new tasks. Our empirical conclusions are consistently supported by experiments on simulated manipulation tasks, and by 52 unique fine-tuning experiments on a real robotic grasping system pre-trained on 580,000 grasps. For video results and an overview of the methods and experiments in this study, see the project website at <https://sites.google.com/view/adapt-on-the-fly-vlrr2020>.

I. INTRODUCTION

The ability to constantly learn, adapt, and evolve is arguably one of the most important properties of an intelligent agent prepared to exist in the real world. Similarly, our robots should be able to continuously learn and adapt throughout their lifetime to the ever-changing environments that they are deployed in. This is a widely recognized requirement. In fact, there is an entire academic sub-field of lifelong learning [65] that is interested in the problem of agents that never stop learning. Despite the wide interest in this ability, most of the intelligent agents deployed today are not tested for their adaptation capabilities. Even though techniques such as reinforcement learning theoretically provide the ability to perpetually learn from trial and error, this is not how they

are typically evaluated. Instead, the predominant method of acquiring a new task with reinforcement learning is to initialize a policy from scratch, collect entirely new data in a stationary environment, and evaluate a static policy that was trained with this data.

This static paradigm does not evaluate the robot’s capability to adapt. It also traps robotic reinforcement learning in the worst-case regime for sample efficiency: the cost to acquire a new task is dominated by sample efficiency of the learning algorithm and the complexity of the task, as reflected in cost of acquiring diverse task data starting from naïve (e.g. random) exploration.

Most machine learning models successfully deployed in the real world, such as those used for computer vision and natural language processing (NLP) do not live in this regime. For instance, the predominant method of acquiring a new computer vision task is to start learning the new task with a pre-trained model for a related task, acquired from a pre-collected data set, and *fine-tune* that model to achieve the new task [11, 24, 10]. This changes the sample efficiency regime of the learning process from one which is dominated by *task complexity* to one that is dominated by *task novelty*, i.e. the difference between the new task and the task on which the model was pre-trained. While a number of works have studied how to use pre-trained ImageNet [9] features for robotics [68, 25, 34], there are remarkably few works that study how to adapt motor skills themselves. Our work attempts to bridge this gap.

We adapt an image-based grasping policy to changes in background, object shape and appearance, lighting conditions, and robot morphology and kinematics, while using less than 0.2% of the data necessary to learn the same task from scratch (see Fig. 1). Our results, supported by simulation and extensive real-world experiments, indicate that a pre-adaptation policy acquired for a task using reinforcement learning can be used to acquire policies for nearby tasks using very little new data and a simple update procedure. Furthermore, we find that this approach of adapting pre-trained policies with off-policy

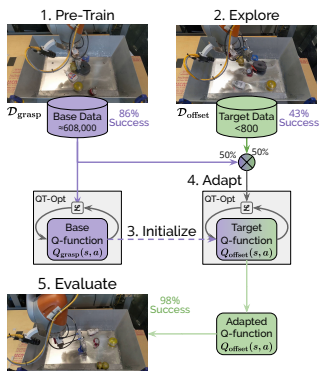


Fig. 2: Schematic of the simple method we test in Section III, using the conceptual framework we discuss in Appendix B-A. We pre-train a policy using the old data from the pre-training task, which is then adapted using the new data from the fine-tuning task.

reinforcement learning (RL) leads to substantial improvements over the course of fine-tuning, and that pre-training via RL is essential: it significantly outperforms conventional pre-training techniques using supervised learning on task-agnostic datasets. We believe this simple adaptation scheme provides a promising solution for creating a lifelong learning robotic agent, and show this potential using a simple continual learning experiment.

To our knowledge, this work is the first to demonstrate that simple fine-tuning of off-policy reinforcement learning can successfully adapt to substantial task, robot, and environment variations which were not present in the original training distribution (*i.e.* off-distribution).

II. RELATED WORK

We consider how we might transfer knowledge for efficient learning in new conditions [63, 46, 61], a widely-studied problem particularly outside of the robotics domain [11, 24, 10, 7, 51]. Prior works in robotics have considered how we might transfer information from models trained with supervised learning on ImageNet [9] by fine-tuning [36, 13, 17, 49] or other means [60, 20]. Our experiments show that transfer from pre-trained conditions is significantly more successful than transfer from ImageNet. Other works have leveraged experience in simulation [56, 66, 57, 62, 45, 55, 48, 22, 19] or representations learned with auxiliary losses [53, 39, 58] for effective transfer. While successful, these approaches either require significant engineering effort to construct an appropriate simulation or significant supervision. Most relevantly, recent work in model-based RL has used predictive models for fast transfer to new experimental set-ups [4, 18], *i.e.* by fine-tuning predictive models [8], via online search of a pre-learned representation of the space models, policies, or high-level skills [5, 6, 30, 38], or by learning physics simulation parameters from real data [52, 28]. We show how fine-tuning is successful with a model-free RL approach, and show how a state-of-the-art grasping system can be adapted to new conditions.

III. A VERY SIMPLE FINE-TUNING METHOD

We define then evaluate a simple technique for offline fine-tuning.

Our experiments model an “on the job” adaptation scenario, where a robot is initially trained to perform a general task (in our case, grasping diverse objects), and then the conditions of the task change in a drastic and substantial way as the robot performs the task, *e.g.* through the introduction of significantly brighter lighting, or a peculiar and unexpected type of object. The robot must adapt to this change quickly in order to recover a proficient policy. Handling these changes reflects what we expect to be a common requirement of reinforcement learning policies deployed in the real world: since an RL policy can learn from all of the experience that it has collected, there is no need to separate learning into clearly distinct training and deployment phases. Instead, it is likely desirable to allow the policy to simply continue learning “on the job” so as to adapt to these changes.

A. The Method

We define a very simple fine-tuning procedure for off-policy RL, as follows (Fig. 2).

First, we (1) pre-train a general grasping policy, as describe in Appendix A-A and [29]. To fine-tune a policy onto a new target task, we (2) use the pre-trained policy to collect an exploration dataset of attempts on the target task; then (3) initialize the same off-policy reinforcement learning algorithm which was used for pre-training (QT-Opt, in our case) with the parameters of the pre-trained policy, and both the target task and base task datasets¹ as the data sources (*e.g.* replay buffers); we then (4) update the policy with this training algorithm, using a reduced learning rate, and sampling training examples with equal probability from the base and target task datasets, for some number of update steps. Finally, we (5) evaluate the fine-tuned policy on the target task.

Our method is offline, *i.e.* it uses a single dataset of target task attempts, and requires no robot interaction after initial dataset collection to compute a fine-tuned policy, which may then be deployed onto a robot.

B. Evaluating offline fine-tuning for real-world grasping

We now turn our attention to how to evaluating this simple method’s effectiveness as an adaptation procedure for end-to-end robot learning, and perhaps continual learning. Our goal is to determine whether the method is sample efficient, whether it works over a broad range of possible variations, and to determine whether it performs better than simpler ways of acquiring the target tasks.

With this goal in mind, we conduct a large panel of ablation experiments experiments on a real 7 DoF Kuka arm. These experiments evaluate the performance of our method across the diverse range of Challenge Tasks (See Appendix A) and a continuum of target task dataset sizes, and compare this performance to two comparison methods.

¹We assume this dataset was saved during training of the base policy

Challenge Task	Original Policy	Ours (exploration grasps)							Comparisons	
		25	50	100	200	400	800	Best (Δ)	Scratch	ImageNet
Checkerboard Backing	50%	67%	48%	71%	47%	89%	90%	90% (+40)	0%	0%
Harsh Lighting	32%	23%	16%	52%	44%	58%	63%	63% (+31)	4%	2%
Extend Gripper 1 cm	75%	93%	67%	80%	51%	90%	69%	93% (+18)	0%	14%
Offset Gripper 10 cm	43%	73%	50%	60%	56%	91%	98%	98% (+55)	37%	47%
Transparent Bottles	49%	46%	43%	65%	65%	58%	66%	66% (+17)	27%	20%
Baseline Grasping Task	86%	98%	81%	84%	78%	93%	89%	98% (+12)	0%	12%

TABLE I: Summary of grasping success rates ($N \geq 50$) for the experiments by challenge task, fine-tuning method, and number of exploration grasps. The experiments “Scratch” and “ResNet 50 + ImageNet” both use 800 exploration grasps and the same update process as the other experiments. “Scratch” starts the grasping network with randomly-initialized parameters. “ResNet 50 + ImageNet” refers to training a grasping network with an equivalent architecture to the other experiments, but with its convolutional layers replaced with a ResNet 50 architecture and pre-loaded with ImageNet features; the non-CNN parts of the network (MLPs for the action inputs and the Q-value output) are randomly-initialized.

The experiments are very challenging. The Transparent Bottles task in particular presents a major challenge to most grasping systems: the transparent bottles generally confuse depth-based sensors and, especially in cluttered bins, require the robot to singulate individual items and position the gripper in the right orientation for grasping. Although our base policy uses only RGB images, it is still not able to grasp the glass bottles reliably, because they differ so much from the objects it observed during training. However, after fine-tuning with only 1 hour (100 grasp attempts) of experience, we observe that the transparent bottles can be picked up with a success rate of 66%, 20% better than the base policy. Figure 4 shows how the robot’s view changes for each challenge task. Note the extreme glare and robot reflections visible in images from the Harsh Lighting challenge.

For videos of our experimental results, see the project website.²

a) Collect datasets: First, we collect a dataset of 800 grasp attempts for each of our 5 challenge tasks (see Table III) plus the base grasping task. We then partitioned each dataset into 6 tiers of difficulty by number of exploration grasps (25, 50, 100, 200, 400, and 800 grasp attempts), yielding 36 individual datasets.

b) Train fine-tuned policies: We train a fine-tuned policy for each of these 36 datasets using the procedure described above. We execute the fine-tuning algorithm for 500,000 gradient steps (see Appendix C for more information on how we chose this number) and use a learning rate of 10^{-4} , which is 25% of learning rate used for pre-training. This yields 36 fine-tuned policies, each trained with a different combination of target task and target dataset size. This set of 36 policies includes 6 policies fine-tuned on data from the base grasping task, for validation.

c) Train comparisons: To provide points of comparison, we train two additional policies for each challenge task and the base grasping task, yielding 12 additional policies.

The first comparison (“Scratch”) is a policy trained using the aforementioned fine-tuning procedure and an 800-grasp data set, but using a randomly-initialized Q-function rather than the Q-function obtained from pre-training. The purpose of this comparison is to help us assess the contribution of the pre-trained parameters to the fine-tuning process’ performance.

The second comparison (“ImageNet”) is also trained using an identical fine-tuning procedure and the 800-grasp dataset, but uses a modified Q-function architecture in which we replace the convolutional trunk of the network with that of the popular ResNet50 architecture [21], initialized with the weights obtained by training the network to classify images from the ImageNet dataset [9]. Refer to Fig. 9 for a diagram of the unmodified architecture. We initialize the remaining fully-connected layers with random parameters, and concatenate the action input features at the end of the CNN (rather than the adding them in middle of the CNN, as in the original architecture). Note that in this comparison, the fine-tuning process still updates all parameters, including those of the ResNet50 sub-network. The purpose of this comparison is to provide a comparison to a strong alternative to end-to-end RL for obtaining pre-training parameters.

d) Evaluate performance: Finally, we evaluate all 48 policies on their target task by deploying them to the robot and executing 50 or more grasp attempts to calculate the policy’s final performance. To reduce the variance of our evaluation statistics, we shuffle the contents of the bin between each trial by executing a randomly-generated sequence of sweeping movements with the end-effector.

The full experiment required more than 15,000 grasp attempts and 14 days of real robot time, and was conducted over approximately one month.

We present a full summary of our results in Table I. Across the board, we observe substantial benefits arising from fine-tuning, suggesting that the robot can indeed adapt to drastically new condition with a modest amount of data: our most data-intensive experiment uses just 0.2% of the data used train the base grasping policy to similar performance. Our method consistently outperforms both the “ImageNet” and “Scratch” comparison methods. We provide more detailed analysis of this experiment in the next section.

The experiments are very challenging. For example, the “Transparent Bottles” task presents a major challenge to most grasping systems: the transparent bottles generally confuse depth-based sensors and, especially in cluttered bins, require the robot to singulate individual items and position the gripper in the right orientation for grasping. Although our base policy uses only RGB images, it is still not able to grasp the transparent bottles reliably, because they differ so much from

²For video results, see <https://ryanjulian.me/continual-fine-tuning>

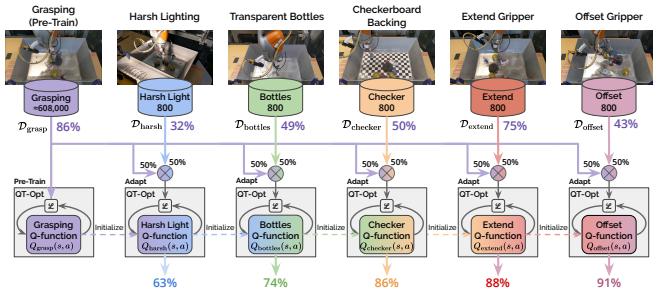


Fig. 3: Flow chart of the continual learning experiment, in which we fine-tune on a sequence of conditions. Every transition to a new scenario happens after 800 grasps.

the objects it observed during training. However, after fine-tuning with only 1 hour (100 grasp attempts) of experience, we observe that the transparent bottles can be picked up with a success rate of 66%, 20% better than the base policy. Similarly, the “Checkerboard Backing” challenge task asks the robot to differentiate edges associated with real objects from edges on an adversarial checkerboard pattern. It never needed this capability to succeed during pre-training, where the background is always featureless and grey, and all edges can be assumed to be associated with a graspable object. After 1 hour (100 grasp attempts) of experience, using our method the robot can grasp objects on the checkerboard background with a 71% success rate, 21% better than the base policy, and this success rate reaches 90% after 8 hours of experience (800 grasp attempts).

C. Evaluating Offline Fine-Tuning for Continual Learning

Now that we have defined and evaluated a simple method for offline fine-tuning, we evaluate its suitability for use in continual learning, which could allow us to achieve the goal of a robot which adapts to ever-changing environments and tasks. To do so, we define a simple continual learning challenge as follows (Fig. 3).

As in the fine-tuning experiments, we begin with a base policy pre-trained for general object grasping. Likewise, we also use our fine-tuning method to adapt the base policy to a target task, in this case “Harsh Lighting.” Not content to stop there, we use *this* adapted policy—not the base policy—as the initialization for another iteration of our fine-tuning algorithm, this time targeting “Transparent Bottles.” We repeat this process until we have run out of new tasks, ending at the task “Offset Gripper 10cm,” at which point we evaluate the policy on the last task.

We perform this experiment using 800 exploration-grasp datasets for each Challenge Task from our ablation study of online fine-tuning with real robots. We summarize the results in Table II. Note that because it is the first step of the continual learning experiment, the policy for “Harsh Lighting” is identical to that of the 800-grasp variant of the single-step experiment.

Recall that our goal for this experiment is to determine whether continual fine-tuning incurs a significant performance penalty compared to the single-step variant, because we are

Challenge Task	Continual Learning	Δ	
		Base	Single
Harsh Lighting	63%	+32%	-
Transparent Bottles	74%	+25%	+8%
Checkerboard Backing	86%	+36%	-4%
Extend Gripper 1 cm	88%	+12%	-5%
Offset Gripper 10 cm	91%	+44%	-7%

TABLE II: Summary of grasping success rates ($N \geq 50$) for the continual learning experiment by challenge task, and comparison to single-step fine-tuning. “Base” refers to the baseline grasping policy before fine-tuning, and “Single” refers to the best performance from the single-step fine-tuning experiment in Table I. Note that because it is the first step of the continual learning experiment, the policy for “Harsh Lighting” is identical to that of the 800-grasp variant of the single-step experiment.

interested in using this method as a building block for continual learning algorithms. We find that continual fine-tuning does not impose a drastic performance penalty compared to single-step fine-tuning. The continual fine-tuning policies for the “Checkerboard Backing,” “Extend Gripper 1 cm,” and “Offset Gripper 10 cm,” challenges succeeded in grasping between 4% and 7% less often than their single-step fine-tuning counterparts, whereas the policy for the challenging “Transparent Bottles” case actually succeeded 8% more often. These small deltas are within the margin-of-error of our evaluation procedure, so we conclude that the effect of continual fine-tuning on the performance compared to single-step fine-tuning is very small. This experiment demonstrates that our method can perform continual adaptation, and may serve as the basis for a continual end-to-end robot learning method.

IV. CONCLUSION

For robots to be able to operate in unconstrained environments, they must be able to continuously adapt to new situations. Our large-scale study shows that combining off-policy RL with a very simple fine-tuning procedure is an effective adaptation method, and this method is capable of achieving remarkable improvements in robot performance on new tasks with very little new data. Furthermore, our continual learning experiment shows that using this simple method in a continual setting imposes very little performance penalty compared to the single-step setting. This suggests that the combination of off-policy RL and fine-tuning can serve as a building block for future continual learning methods. Our results comparing supervised-learning-based initialization to those acquired with our RL-fine-tuning approach highlight a familiar truism about robotics: that robotic agents must do more than perceive the world, they must also act in it. The ability to learn the combination of these two capabilities is what makes RL well-suited for creating continually-learning robots.

ACKNOWLEDGMENTS

The authors thank Noah Brown and Ivonne Fajardo for their superb and unyielding support with real robot experiments. We also thank Alex Irpan and Eric Jang for their help with robot learning software, Yevgen Chebotar for his advice on

early revisions of this work and always-insightful discussions, Dmitry Kalashnikov and Jake Varley for their help with QT-Opt, and K.R. Zentner for her help with editing and artwork for this paper.

REFERENCES

- [1] Pulkit Agrawal, Ashvin V Nair, Pieter Abbeel, Jitendra Malik, and Sergey Levine. Learning to poke by poking: Experiential learning of intuitive physics. In *Advances in neural information processing systems*, pages 5074–5082, 2016.
- [2] Ferran Alet, Tomás Lozano-Pérez, and Leslie P Kaelbling. Modular meta-learning. *arXiv preprint arXiv:1806.10166*, 2018.
- [3] Alessandro Bonardi, Stephen James, and Andrew J Davison. Learning one-shot imitation from humans without humans. *arXiv preprint arXiv:1911.01103*, 2019.
- [4] Konstantinos Chatzilygeroudis and Jean-Baptiste Mouret. Using parameterized black-box priors to scale up model-based policy search for robotics. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–9. IEEE, 2018.
- [5] Konstantinos Chatzilygeroudis, Vassilis Vassiliades, and Jean-Baptiste Mouret. Reset-free trial-and-error learning for robot damage recovery. *Robotics and Autonomous Systems*, 100:236–250, 2018.
- [6] Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. Robots that can adapt like animals. *Nature*, 521(7553):503–507, 2015.
- [7] Wenyuan Dai, Qiang Yang, Gui-Rong Xue, and Yong Yu. Boosting for transfer learning. In *Proceedings of the 24th international conference on Machine learning*, pages 193–200, 2007.
- [8] Sudeep Dasari, Frederik Ebert, Stephen Tian, Suraj Nair, Bernadette Bucher, Karl Schmeckpeper, Siddharth Singh, Sergey Levine, and Chelsea Finn. Robonet: Large-scale multi-robot learning, 2019.
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [11] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *International conference on machine learning*, pages 647–655, 2014.
- [12] Chelsea Finn and Sergey Levine. Deep visual foresight for planning robot motion. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2786–2793. IEEE, 2017.
- [13] Chelsea Finn, Xin Yu Tan, Yan Duan, Trevor Darrell, Sergey Levine, and Pieter Abbeel. Deep spatial autoencoders for visuomotor learning. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 512–519. IEEE, 2016.
- [14] Chelsea Finn, Tianhe Yu, Tianhao Zhang, Pieter Abbeel, and Sergey Levine. One-shot visual imitation learning via meta-learning. *arXiv preprint arXiv:1709.04905*, 2017.
- [15] Ali Ghadirzadeh, Atsuto Maki, Danica Kragic, and Mårten Björkman. Deep predictive policy training using reinforcement learning. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2351–2358. IEEE, 2017.
- [16] Vijaykumar Gullapalli, Judy A Franklin, and Hamid Benbrahim. Acquiring robot skills via reinforcement learning. *IEEE Control Systems Magazine*, 14(1):13–24, 1994.
- [17] Abhinav Gupta, Adithyavairavan Murali, Dhiraj Gandhi, and Lerrel Pinto. Robot learning in homes: Improving generalization and reducing dataset bias, 2018.
- [18] David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. In *Advances in Neural Information Processing Systems*, pages 2450–2462, 2018.
- [19] Aleksi Hämmäläinen, Karol Arndt, Ali Ghadirzadeh, and Ville Kyrki. Affordance learning for end-to-end visuomotor robot control. *arXiv preprint arXiv:1903.04053*, 2019.
- [20] Murtaza Hazara and Ville Kyrki. Transferring generalizable motor primitives from simulation to real world. *IEEE Robotics and Automation Letters*, 4(2):2172–2179, 2019.
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [22] Juan Camilo Gamboa Higuera, David Meger, and Gregory Dudek. Adapting learned robotics behaviours through policy adjustment. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5837–5843. IEEE, 2017.
- [23] Jeremy Howard and Sebastian Ruder. Universal language model fine-tuning for text classification. *arXiv preprint arXiv:1801.06146*, 2018.
- [24] Jeremy Howard and Sebastian Ruder. Universal language model fine-tuning for text classification, 2018.
- [25] Minyoung Huh, Pulkit Agrawal, and Alexei A Efros. What makes imagenet good for transfer learning? *arXiv preprint arXiv:1608.08614*, 2016.
- [26] Alexander Irpan, Kanishka Rao, Konstantinos Bousmalis, Chris Harris, Julian Ibarz, and Sergey Levine. Off-policy evaluation via off-policy classification. In *Advances in Neural Information Processing Systems*, pages 5438–5449, 2019.
- [27] Stephen James, Michael Bloesch, and Andrew J Davison. Task-embedded control networks for few-shot imitation learning. In *Conference on Robot Learning*, pages 783–795, 2018.

- [28] Rae Jeong, Jackie Kay, Francesco Romano, Thomas Lampe, Tom Rothorl, Abbas Abdolmaleki, Tom Erez, Yuval Tassa, and Francesco Nori. Modelling generalized forces with reinforcement learning for sim-to-real transfer. *arXiv preprint arXiv:1910.09471*, 2019.
- [29] Dmitry Kalashnikov, Alex Irpan, Peter Pastor, Julian Ibarz, Alexander Herzog, Eric Jang, Deirdre Quillen, Ethan Holly, Mrinal Kalakrishnan, Vincent Vanhoucke, et al. Scalable deep reinforcement learning for vision-based robotic manipulation. In *Conference on Robot Learning*, pages 651–673, 2018.
- [30] Rituraj Kaushik, Pierre Desreumaux, and Jean-Baptiste Mouret. Adaptive prior selection for repertoire-based online adaptation in robotics. *Frontiers in Robotics and AI*, 6:151, 2020.
- [31] Jens Kober and Jan R Peters. Policy search for motor primitives in robotics. In *Advances in neural information processing systems*, pages 849–856, 2009.
- [32] Jens Kober, J Andrew Bagnell, and Jan Peters. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11):1238–1274, 2013.
- [33] Nate Kohl and Peter Stone. Machine learning for fast quadrupedal locomotion. In *The Nineteenth National Conference on Artificial Intelligence*, pages 611–616, July 2004.
- [34] Simon Kornblith, Jonathon Shlens, and Quoc V Le. Do better imagenet models transfer better? In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2661–2671, 2019.
- [35] Michelle A Lee, Yuke Zhu, Krishnan Srinivasan, Parth Shah, Silvio Savarese, Li Fei-Fei, Animesh Garg, and Jeannette Bohg. Making sense of vision and touch: Self-supervised learning of multimodal representations for contact-rich tasks. *arXiv preprint arXiv:1810.10191*, 2018.
- [36] Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research*, 17(1):1334–1373, 2016.
- [37] Sridhar Mahadevan and Jonathan Connell. Automatic programming of behavior-based robots using reinforcement learning. *Artificial intelligence*, 55(2-3):311–365, 1992.
- [38] Josh Merel, Saran Tunyasuvunakool, Arun Ahuja, Yuval Tassa, Leonard Hasenclever, Vu Pham, Tom Erez, Greg Wayne, and Nicolas Heess. Reusable neural skill embeddings for vision-guided whole body movement and object manipulation. *arXiv preprint arXiv:1911.06636*, 2019.
- [39] Piotr Mirowski, Razvan Pascanu, Fabio Viola, Hubert Soyer, Andrew J Ballard, Andrea Banino, Misha Denil, Ross Goroshin, Laurent Sifre, Koray Kavukcuoglu, et al. Learning to navigate in complex environments. *arXiv preprint arXiv:1611.03673*, 2016.
- [40] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- [41] Anusha Nagabandi, Ignasi Clavera, Simin Liu, Ronald S Fearing, Pieter Abbeel, Sergey Levine, and Chelsea Finn. Learning to adapt in dynamic, real-world environments through meta-reinforcement learning. 2018.
- [42] Anusha Nagabandi, Chelsea Finn, and Sergey Levine. Deep online learning via meta-learning: Continual adaptation for model-based rl. *arXiv preprint arXiv:1812.07671*, 2018.
- [43] Anusha Nagabandi, Kurt Konoglie, Sergey Levine, and Vikash Kumar. Deep dynamics models for learning dexterous manipulation. *arXiv preprint arXiv:1909.11652*, 2019.
- [44] Ashvin V Nair, Vitchyr Pong, Murtaza Dalal, Shikhar Bahl, Steven Lin, and Sergey Levine. Visual reinforcement learning with imagined goals. In *Advances in Neural Information Processing Systems*, pages 9191–9200, 2018.
- [45] OpenAI, Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, Jonas Schneider, Nikolas Tezak, Jerry Tworek, Peter Welinder, Lilian Weng, Qiming Yuan, Wojciech Zaremba, and Lei Zhang. Solving rubik’s cube with a robot hand, 2019.
- [46] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2009.
- [47] Deepak Pathak, Parsa Mahmoudieh, Guanghao Luo, Pulkit Agrawal, Dian Chen, Fred Shentu, Evan Shelhamer, Jitendra Malik, Alexei A. Efros, and Trevor Darrell. Zero-shot visual imitation. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, Jun 2018. doi: 10.1109/cvprw.2018.00278. URL <http://dx.doi.org/10.1109/CVPRW.2018.00278>.
- [48] Xue Bin Peng, Marcin Andrychowicz, Wojciech Zaremba, and Pieter Abbeel. Sim-to-real transfer of robotic control with dynamics randomization. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 1–8. IEEE, 2018.
- [49] Lerrel Pinto and Abhinav Gupta. Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours. In *2016 IEEE international conference on robotics and automation (ICRA)*, pages 3406–3413. IEEE, 2016.
- [50] Vitchyr H Pong, Murtaza Dalal, Steven Lin, Ashvin Nair, Shikhar Bahl, and Sergey Levine. Skew-fit: State-covering self-supervised reinforcement learning. *arXiv preprint arXiv:1903.03698*, 2019.
- [51] Rajat Raina, Alexis Battle, Honglak Lee, Benjamin Packer, and Andrew Y Ng. Self-taught learning: transfer learning from unlabeled data. In *Proceedings of the 24th international conference on Machine learning*, pages

- 759–766, 2007.
- [52] Divyam Rastogi, Ivan Koryakovskiy, and Jens Kober. Sample-efficient reinforcement learning via difference models.
- [53] Martin A. Riedmiller, Roland Hafner, Thomas Lampe, Michael Neunert, Jonas Degraeve, Tom Van de Wiele, Volodymyr Mnih, Nicolas Manfred Otto Heess, and Jost Tobias Springenberg. Learning by playing solving sparse reward tasks from scratch. In *ICML*, 2018.
- [54] Sebastian Ruder. An overview of multi-task learning in deep neural networks, 2017.
- [55] Andrei A. Rusu, Matej Vecerík, Thomas Rothörl, Nicolas Manfred Otto Heess, Razvan Pascanu, and Raia Hadsell. Sim-to-real robot learning from pixels with progressive nets. In *CoRL*, 2016.
- [56] Fereshteh Sadeghi and Sergey Levine. Cad2rl: Real single-image flight without a single real image. *Robotics: Science and Systems XIII*, Jul 2017. doi: 10.15607/rss.2017.xiii.034. URL <http://dx.doi.org/10.15607/RSS.2017.XIII.034>.
- [57] Fereshteh Sadeghi, Alexander Toshev, Eric Jang, and Sergey Levine. Sim2real viewpoint invariant visual servoing by recurrent control. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 4691–4699, 2018. doi: 10.1109/CVPR.2018.00493. URL http://openaccess.thecvf.com/content_cvpr_2018/html/Sadeghi_Sim2Real_Viewpoint_Invariant_CVPR_2018_paper.html.
- [58] Alexander Sax, Bradley Emi, Amir R. Zamir, Leonidas J. Guibas, Silvio Savarese, and Jitendra Malik. Mid-level visual representations improve generalization and sample efficiency for learning visuomotor policies. In *Conference on Robot Learning*, 2019.
- [59] Gerrit Schoettler, Ashvin Nair, Jianlan Luo, Shikhar Bahl, Juan Aparicio Ojea, Eugen Solowjow, and Sergey Levine. Deep reinforcement learning for industrial insertion tasks with visual inputs and natural reward signals. In *International Conference on Learning Representations*, 2019.
- [60] Pierre Sermanet, Kelvin Xu, and Sergey Levine. Unsupervised perceptual rewards for imitation learning. *Proceedings of Robotics: Science and Systems (RSS)*, 2017. URL <http://arxiv.org/abs/1612.06699>.
- [61] Chuanqi Tan, Fuchun Sun, Tao Kong, Wenchang Zhang, Chao Yang, and Chunfang Liu. A survey on deep transfer learning. *Lecture Notes in Computer Science*, page 270279, 2018. ISSN 1611-3349. doi: 10.1007/978-3-030-01424-7_27. URL http://dx.doi.org/10.1007/978-3-030-01424-7_27.
- [62] Jie Tan, Tingnan Zhang, Erwin Coumans, Atil Iscen, Yunfei Bai, Danijar Hafner, Steven Bohez, and Vincent Vanhoucke. Sim-to-real: Learning agile locomotion for quadruped robots. *Robotics: Science and Systems XIV*, Jun 2018. doi: 10.15607/rss.2018.xiv.010. URL <http://dx.doi.org/10.15607/RSS.2018.XIV.010>.
- [63] Matthew E Taylor and Peter Stone. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10(Jul):1633–1685, 2009.
- [64] Russell L. Tedrake. *Applied optimal control for dynamically stable legged locomotion*. Thesis, Massachusetts Institute of Technology, 2004. URL <https://dspace.mit.edu/handle/1721.1/28742>.
- [65] Sebastian Thrun. *Lifelong Learning Algorithms*, page 181209. Kluwer Academic Publishers, USA, 1998. ISBN 0792380479.
- [66] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep 2017. doi: 10.1109/iros.2017.8202133. URL <http://dx.doi.org/10.1109/IROS.2017.8202133>.
- [67] Lin Yen-Chen, Maria Bauza, and Phillip Isola. Experience-embedded visual foresight. *arXiv preprint arXiv:1911.05071*, 2019.
- [68] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? In *Advances in neural information processing systems*, pages 3320–3328, 2014.
- [69] Tianhe Yu, Chelsea Finn, Annie Xie, Sudeep Dasari, Tianhao Zhang, Pieter Abbeel, and Sergey Levine. One-shot imitation from observing humans via domain-adaptive meta-learning. In *International Conference on Learning Representations*, 2018.
- [70] Tianhe Yu, Gleb Shevchuk, Dorsa Sadigh, and Chelsea Finn. Unsupervised visuomotor control through distributional planning networks. *arXiv preprint arXiv:1902.05542*, 2019.
- [71] Andy Zeng, Shuran Song, Stefan Welker, Johnny Lee, Alberto Rodriguez, and Thomas Funkhouser. Learning synergies between pushing and grasping with self-supervised deep reinforcement learning. *arXiv preprint arXiv:1803.09956*, 2018.
- [72] Andy Zeng, Shuran Song, Johnny Lee, Alberto Rodriguez, and Thomas Funkhouser. Tossingbot: Learning to throw arbitrary objects with residual physics. *arXiv preprint arXiv:1903.11239*, 2019.

APPENDIX A

THE ROBUSTNESS OF LEARNED POLICIES: A CASE STUDY

To study the problem of adaptation, we utilize a grasping policy pre-trained with RL, which we evaluate in five different conditions that were not encountered during pre-training. In this section, we will describe the pre-training process and test the robustness of the pre-trained policy to various robot and environment modifications. We choose these modifications to reflect changes we believe a learning robot would experience, and should be expected to adapt to, when deployed “on the job” in the real world. In Section III, we will describe a simple fine-tuning based adaptation process, and evaluate it using these modifications.

Challenge Task	Type	Base Policy	Δ
Checkerboard Backing	Background	50%	-36%
Harsh Lighting	Lighting conditions	31%	-55%
Extend Gripper 1 cm	Gripper shape	76%	-10%
Offset Gripper 10 cm	Robot morphology	47%	-39%
Transparent Bottles	Unseen objects	49%	-37%

TABLE III: Summary of modifications to the robot and environment, and their effect on the performance of the base policy. Changing the background lighting, morphology, and objects leads to substantial degradation in performance compared to the original training conditions.

A. Pre-training process

We pre-train the grasping policy, which we refer to as the “base policy,” using the QT-Opt algorithm in two stages, as described in [29]. First, we train a Q-function network offline using data from 580,000 real grasp attempts over a corpus of 1,000 visually and physically diverse objects. Second, we continue training this network online³ over the course of 28,000 real grasp attempts on the same corpus of objects. That is, we use a real robot to collect trials using the current network, update the network using these new trials, deploy the updated network to the real robot, and repeat. This procedure yields a final base policy that achieves 96% accuracy on a set of previously-unseen test objects. We use a challenging subset of six of these test objects for most experiments in this work. On this set, **our base model achieves a success rate of 86% on the baseline grasping task.**

B. Robustness of the pre-trained policy

We begin by choosing set of significant modifications to the robot and environment, which we believe are characteristic of a real-world continual learning scenario. We then evaluate the performance of the base policy on increasingly-severe versions of these modifications. This process allows us to assess the limits of robustness of policies trained using the pre-training method. Once we find a modification that is sufficiently-severe to compromise the base policy’s performance in each category, we use it to define a “Challenge Task” for our study of adaptation methods.

Next, we describe these challenges and the corresponding performance of the base policy.

Background: We introduce a black-white 1 inch checkerboard pattern that we glue to the bottom of the robot’s workspace (see Fig. 1, fourth from left). We observe that conventional variations in the workspace surface, such as uniform changes in color or specularity, have no effect on the base policy’s performance. Introducing a checkerboard pattern often fools the robot into grasping at checkerboard edges rather than objects. This adversarial modification compromises the base policy’s performance to 50% (-36% compared to the base task).

³Following the example set by [29], we refer this procedure as “online” rather than “on-policy,” because the policy is still updated by the off-policy reinforcement learning algorithm




Fig. 4: Views of from the robot camera for each of our six Challenge Tasks and the base grasping task.

Lighting conditions: We introduce a high-intensity halogen light source parallel with the workspace (see Fig. 1, second from left), creating a bright spot in the robot’s camera view, and intense light-dark contrasts along the plane of the workspace. The base policy was trained in standard indoor lighting conditions, with no exposure to natural light or significant variation. We observe that mild perturbations in lighting conditions (i.e. those which can be created by standard-intensity household lights) have no effect on the base policy’s performance. Using the very bright halogen light source has a severe impact, and degrades the base policy’s performance to 31% (-55% compared to the baseline).

Gripper shape: We extend the parallel gripper attached to the robot by 1 cm and significantly narrow its width and compliance in the process (see Fig. 1, fifth from left). This changes the robot’s kinematics (lengthening the gripper in the distal direction), while also lowering the relative pose of the robot with respect to the workspace surface by 1 cm. This modification compromises the base policy’s performance to 76% (-10% compared to the baseline).

Robot morphology: We translate the gripper laterally by 10 cm (see Fig. 1, far-right). Note that during training this policy experienced absolutely no variation in robot morphology. We observe that translating the gripper laterally by up to 5 cm has no impact on performance. By translating the gripper laterally by 10 cm (approximately a full gripper or arm link width), we degrade the base policy’s performance to 47% (-39% compared to the baseline).

Unseen objects: We introduce completely-transparent plastic beverage bottles (see Fig. 1, third from left) that were not present in the training set. Based on our experiments, the system is robust to a broad variety of previously-unseen objects, as long as they have significant opaque components. For example, even though there are no drinking bottles in the training set, we find the system is able to pick up labeled drink bottles with 98% success rate. Success rates for other novel, opaque objects are similarly consistent with the baseline performance on the test set. However, we find that introducing completely-transparent drink bottles causes the base policy to

often grasp where two bottles are adjacent, *i.e.* as though it cannot differentiate which parts of the scene are inside vs outside a bottle. By introducing completely-transparent plastic beverage bottles, we are able to compromise the base policy’s performance to 49% (-37% compared to the baseline).

See Table III for a summary of the modification experiments, and their effect on base policy performance.

APPENDIX B

ASSESSING FINE-TUNING TECHNIQUES FOR END-TO-END ROBOTIC REINFORCEMENT LEARNING

We propose a conceptual framework for fine-tuning algorithms, and use simulation experiments to assess the suitability of some algorithm variations for end-to-end robot learning.

“Fine-tuning” refers to a family of transfer learning techniques, in which we seek to acquire a neural network for one task (which we will refer to as the “target” task) by making use of some or all of a network trained on a related task (the “base” task). This is a very common technique for quickly acquiring new tasks in computer vision [11, 25, 34] and natural language processing [23]. As collecting new robot experience data is expensive, our goal is to use as little target task data as possible. In this section, we first describe the general algorithmic sketch for fine-tuning, then enumerate some of the most common fine-tuning techniques. In Sections III and C, we evaluate the suitability of these techniques for end-to-end robot learning.

A. Fine-Tuning: Conceptual Framework

We can organize fine-tuning for end-to-end reinforcement learning into four essential steps (Fig. 2). Different fine-tuning techniques change the details of one of these steps.

- 1) **Pre-training:** *Pre-train a policy to perform some base task, which is related to our target task.* In the experiments in this work, the base task is always indiscriminate object grasping. In computer vision and NLP, this step can often be skipped by making use of one of many pre-trained and publicly-available state-of-the-art vision and language models. We hope for a future in which this is possible in robotics.
- 2) **Exploration:** *Explore in the new target task, to collect data for adaptation.* In principle, in off-policy reinforcement learning any policy may be used for exploration. In our study, and what we believe to be most representative of a real-world continual learning scenario, we always use the pre-trained policy for exploration.
- 3) **Initialization:** *Initialize the policy for the target task using some or all of the weights from the pre-trained policy.* The standard implementation of this step is to start with the entire pre-trained network. Some techniques may choose to use only a subset of the pre-trained network (*e.g.* truncating the last few layers of a CNN).
- 4) **Adaptation:** *Use the exploration data update the initialized policy to perform the new task.* The standard version of this step continues updating the entire initialized policy with the same algorithm and hyperparameters




Fig. 5: Comparison of fine-tuning performance for a policy which uses all base parameters, and a policy which initializes the head parameters from scratch. Re-initializing parameters has a negative effect on sample efficiency for fine-tuning.

as was used for the pre-training process, but with the target task data. There are many variations on this step, including which parts of the network to update, at what learning rate, with what data, with which optimization algorithm, whether to add additional network layers, etc.

- 5) **Evaluation:** *Assess performance of the fine-tuned network on the new task.* If this step only happens once, we refer to such a technique as “offline fine-tuning,” because the adaptation step never uses data from an updated policy. If this step happens repeatedly (*e.g.* exploration and evaluation are one-and-the-same), and its result is used for further adaptation to the same target task, we refer to a technique as “online fine-tuning.” We explore both variations in our experiments.

Using this fine-tuning framework, we consider several variations of fine-tuning, and assess their suitability for end-to-end robotic RL. Notably, we neglect an analysis of pre-training techniques for fine-tuning reinforcement learning (*i.e.* (1)), which has a large and rapidly-growing body of research in the meta- and multi-task RL communities (see Sec. II). Instead, we focus on initialization (2) and adaptation (3). All of our experiments use end-to-end off-policy reinforcement learning of an indiscriminate object grasping task for their pre-training step. Refer to Section A-A for details on our pre-training process.

B. Experiments in simulation

We use simulation experiments to evaluate the suitability of some fine-tuning variations, along the axes we defined in Section B-A.

- 1) **Adding a new head and other selective initialization techniques:** Selective-initialization techniques start the fine-tuning process with a policy which has some of its parameters initialized to random, *e.g.* a popular variant is to “add a head” to a pre-trained neural network by omitting its last few layer(s)




Fig. 6: Performance curve for an online fine-tuning simulation experiment. The base policy is pre-trained to grasp opaque colored blocks, and the target task is to grasp semi-transparent blocks. Each curve represents a different fraction of target task data, and the remaining data is sampled from the base task. In simulation, the amount of target task data has a straightforward relationship with sample efficiency.

from initialization, so that the new head can be trained to perform on the target task.

Figure 5 portrays a study of partial initialization for online fine-tuning using a simulated grasping experiment. In this experiment, the base task is “grasp opaque blocks” and the target task “grasp semi-transparent blocks,” and the base policy performance is 98% when trained from scratch on 43,000 grasp attempts. Both fine-tuned policies begin with low performance, around 15%. After 5000 exploration grasps (12% of the data used for the base policy), the performance of the full initialization policy has reached the base policy performance, while the policy with a new head has barely reached 30%. This gap shows that the combination of off-policy RL and selective initialization is unsuitable for sample-efficient fine-tuning.

Our experiments immediately make apparent the downsides of selective initialization for fine-tuning. In particular, online fine-tuning requires to maintain a policy that can competently explore the target task at all times, any method which compromises the performance of such a policy—even temporarily—has a high risk of failing as a sample-efficient fine-tuning technique. The resulting performance gap, once created, is hard to recover from. As a consequence, we find in simulation experiments that online fine-tuning with selective re-initialization takes a significant fraction of the pre-training samples to converge to baseline performance, making this family of fine-tuning methods sample inefficient.

C. Training with a mix of data from the base and target tasks

We experiment with mixing data from the pre-training task into the fine-tuning process (Fig. 6), and find that in simulation this has a predictable relationship with sample efficiency:

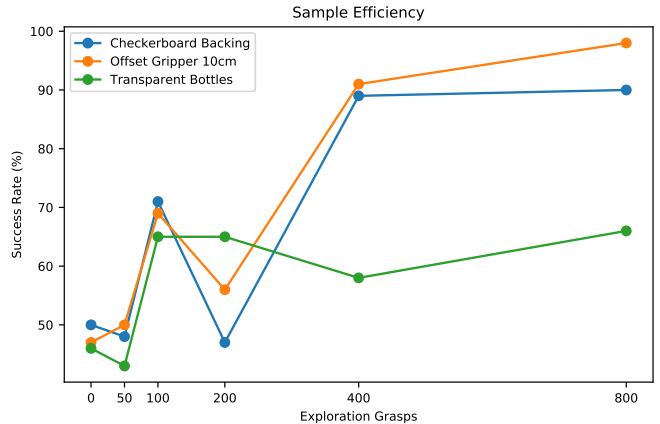


Fig. 7: Sample efficiency of our fine-tuning method on selected real-robot challenge tasks.

higher shares of target task data allow the fine-tuning policy achieve higher performance faster.

Our goal is to design a fine-tuning algorithm for real robots which might be used for continual learning, and our conclusion from this brief study is that online fine-tuning is a poor fit for for this goal. The experiments with selective re-initialization in particular highlights the challenge of online fine-tuning: it only allows us to use algorithms which preserve the exploration ability of the policy at all times. We also believe that offline fine-tuning is more practical than online fine-tuning, due to the inherent complexity of placing a robot in the loop of a reinforcement learning algorithm. If used as part of a continual learning method, an offline method would also allow a robot to collect data on a new task piecemeal, and only attempt to adapt to that new task when it has collected enough data to be successful.

APPENDIX C EMPIRICAL ANALYSIS

In this section, we aim to further investigate the efficiency, performance, and characteristics of our large-scale real-world adaptation experiments.

A. Performance and sample efficiency of our method

Figure 7 shows the success rates for our method from Table I against the amount of data used to achieve that success rate for selected tasks. The data indicates that our simple offline fine-tuning method can adapt policies to many new tasks with performance at or even above the state-of-the-art base policy, using modest amounts of data. For instance, “Extend Gripper 1cm” and “Offset Gripper 10cm” both needed only 25 exploration grasps to achieve substantial gains in performance (+18% and +30%, respectively). All policies attain substantial performance gains over the base policy by the time they are exposed to 800 exploration grasps, which is less than 0.2% of the data necessary to train an equivalently-performing policy on the base task.

While the general trend is that more exploration data leads to higher performance, this relationship is not linear. All




Fig. 8: Evaluation performance of a single offline fine-tuning experiment at different numbers of gradient steps (optimization epochs). The blue curve is real robot performance on the target task (Offset Gripper 10cm) when trained using 400 exploration grasps. The green dotted line is the performance of training the same policy from scratch (random initialization) using 800 exploration grasps, and the yellow dotted line the performance of the base policy. The red dotted line portrays the number of gradient steps we choose to use for our large-scale fine-tuning study.

methods experience a substantial improvement in performance after 100 or fewer exploration grasps. However, we observe that these performance improvements in the very low-data regime (*e.g.* ≤ 200 grasp attempts) are also unstable.

B. The downside of offline fine-tuning: deciding when to stop

Our results indicate that offline fine-tuning can train robotic policies to substantial performance improvements with modest amounts of data, and that offline methods are not limited by the need to preserve an always-sufficient exploration policy as with online methods. However, we identify one significant drawback to the method compared to online fine-tuning.

A pure offline fine-tuning method has no built-in evaluation step which would inform us when the robot’s performance on the target task has stopped improving, and therefore when we should stop fine-tuning with a fixed set of target task data. This is a subset of the off-policy evaluation problem [26]. Knowing when the policy stops improving is important, because fine-tuning exists in a low-data regime, and repeatedly updating a neural network model with small amounts of data leads to overfitting onto that data. Not only does this degrade the performance on the target task, but also the ability of the network to adapt to new tasks later (*i.e.* for continual learning).

We can see this phenomenon in Figure 8 showing a real robot’s performance on the “Offset Gripper 10cm” target task at different numbers of steps into an offline fine-tuning process that uses 400 exploration grasps. Performance quickly rises until around 500,000 gradient steps. Past this point, it precipitously drops and never recovers, dropping below even the initial performance of the base policy from which it was

trained, as the initialization is being overwritten by overfitting to the target samples. The point at which overfitting begins is a function of the initialized model, target dataset, learning algorithm, and many other factors, and is not necessarily stable or easily predictable.

For the purposes of our large-scale fine-tuning study, we use this experiment and several others to determine that 500,000 gradient steps was an acceptable choice for the real-world experiments, but the variance in the results in Table I and Figure 7 shows that this choice was not necessarily optimal for all of our tasks and datasets. We believe one practical a solution to this problem of a continual learning robot is to use a mix of offline fine-tuning and online evaluation. The point, at which performance stops improving represents when the training process has exhausted the fine-tuning dataset of new information, and the robot must return to exploring online to continue improving.

C. Comparing initializing with RL to initializing with supervised learning

In order to answer the question whether RL is better suited for creating a continually-learning robotic agent than supervised learning, we compare our results to an ImageNet-pretrained baseline. The ImageNet baseline uses a similar grasping network where its convolutional layers are replaced with ResNet 50 architecture and pre-loaded with ImageNet features. Since the part of the network that process robot’s state and action inputs cannot be initialized using supervised learning, we initialize them randomly. As shown in Table I, the best performing ImageNet-based agent achieves the success rate of 47% on “Offset Gripper 10cm,” which corresponds to 4% improvement over the base policy performance. This result seems to confirm our hypothesis that our RL-based pre-training is crucial for good subsequent fine-tuning. Note that we first attempted to fine-tune these ImageNet-based policies while holding the ImageNet feature layers constant, but this procedure failed to achieve any non-zero success rate. This suggests that, unlike adapting computer vision networks to new visual tasks, adapting end-to-end robot learning to new sensorimotor tasks may require changing the features used to represent the problem, and not just the post-processing of said features.

Figure 9 highlights some of the changes that happen during the RL-based fine-tuning in greater detail. It demonstrates the normalized distance in parameter space of a fine-tuned policy for each of our challenge tasks from its base policy. While it is unsurprising that primarily-visual challenges such as “Checkerboard Backing” and “Harsh Lighting” induce large changes in the parameters of the convolutional parts of the network, we observe that even “Offset Gripper 10cm,” a purely-morphological change to the robot, induces substantial changes to the network’s image-processing parameters (*e.g.* layers conv2-conv7). We attribute this to the successful agent’s need for hand-eye coordination to complete the task: offsetting the gripper not only changes robot morphology, it changes the location of the robot in its own visual field drastically. In order




Fig. 9: Analysis of parameter changes induced by different fine-tuning target tasks. This plot portrays the cosine distance between the parameters of the pre-trained and fine-tuned networks for our 5 fine-tuning target tasks. The bar heights are normalized by the magnitude of parameter changes induced in the Q-function network by fine-tuning the baseline grasping task.

to perform effective visual servoing with a new morphology, both the image and action-processing parts of the network must be updated.

APPENDIX D ADDITIONAL RELATED WORK

Reinforcement learning is a long-standing approach for enabling robots to autonomously acquire skills [32] such as locomotion [33, 64], pushing objects [37, 12], ball-in-cup manipulation [31], peg insertion [16, 36, 59, 35, 71], throwing objects [15, 72], and grasping [49, 29]. We particularly focus on the problem of deep reinforcement learning from raw pixel observations [36], as it allows us to place little restrictions on state representation. A number of works have also considered this problem setting [13, 15, 12, 71, 1, 40]. However, a key challenge with deep RL methods is that they typically learn each skill from scratch, disregarding previously-learned skills. If we hope for robots to generalize to a broad range of real world environments, this approach is not practical.

Other works have aimed to share and transfer knowledge across tasks and conditions by simultaneously learning across multiple goals and tasks [54]. For example, prior works in model-based RL [12, 67, 43] and in goal-conditioned RL [1, 44, 47, 50, 70] have shared data and representations across multiple goals and objects. Along a similar vein, prior work in robotic meta-learning has aimed to learn representations that can be quickly adapted to new dynamics [41, 2, 42] and objects [14, 27, 69, 3]. We consider adaptation to a broad class of changes including dynamics, object classes, and visual observations, including conditions that shift substantially from the training conditions, and do not require the full set of conditions to be represented during the initial training phase.