ASYMMETRIC SELF-PLAY FOR AUTOMATIC GOAL DISCOVERY IN ROBOTIC MANIPULATION

Anonymous authors
Paper under double-blind review

ABSTRACT

We train a single, goal-conditioned policy that can solve many robotic manipulation tasks, including tasks with previously unseen goals and objects. To do so, we rely on asymmetric self-play for goal discovery, where two agents, Alice and Bob, play a game. Alice is asked to propose challenging goals and Bob aims to solve them. We show that this method is able to discover highly diverse and complex goals without any human priors. We further show that Bob can be trained with only sparse rewards, because the interaction between Alice and Bob results in a natural curriculum and Bob can learn from Alice's trajectory when relabeled as a goal-conditioned demonstration. Finally, we show that our method scales, resulting in a single policy that can generalize to many unseen hold-out tasks such as setting a table, stacking blocks, and solving simple puzzles.

1 INTRODUCTION

We are motivated to train a *single* goal-conditioned policy (Kaelbling, 1993) that can solve *any* robotic manipulation task that a human may request in a given environment. In this work, we make progress towards this goal by solving a robotic manipulation problem in a table-top setting where the robot's task is to change the initial configuration of a variable number of objects on a table to match a given goal configuration. This problem is simple in its formulation but likely to challenge a wide variety of cognitive abilities of a robot as objects become diverse and goals become complex.

Motivated by the recent success of deep reinforcement learning for robotics (Levine et al., 2016; Gu et al., 2017; Hwangbo et al., 2019; OpenAI et al., 2019a), we tackle this problem using deep reinforcement learning on a very large training distribution. An open question in this approach is how we can build a training distribution rich enough to achieve generalization to many unseen manipulation tasks. This involves defining both an environment's initial state distribution and a goal distribution. The initial state distribution determines how we sample a set of objects and their configuration at the beginning of an episode, and the goal distribution defines how we sample target states given an initial state. In this work, we focus on a scalable way to define a rich goal distribution.

The research community has started to explore automated ways of defining goal distributions. For example, previous works have explored learning a generative model of goal distributions (Florensa

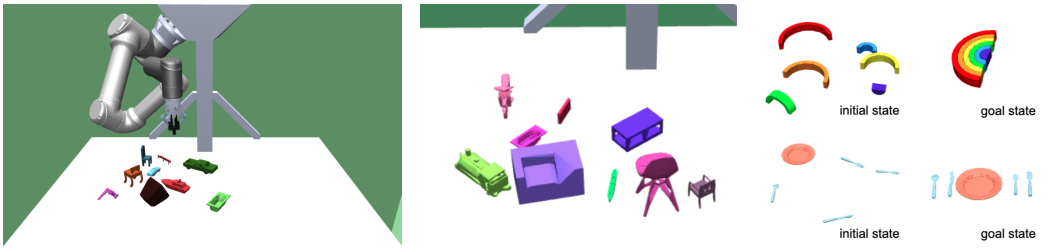

(a) Table-top setting with a robot arm (b) Example initial state for training (c) Example holdout tasks

Figure 1: (a) We train a policy that controls a robot arm operating in a table-top setting. (b) Randomly placed ShapeNet (Chang et al., 2015) objects constitute an initial state distribution for training. (c) We use multiple manually designed holdout tasks to evaluate the learned policy.

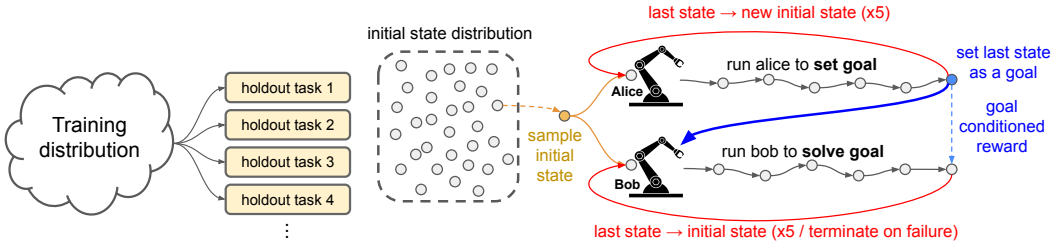

(a) Overall framework (b) Training distribution based on asymmetric self-play

Figure 2: (a) We train a goal-conditioned policy on a single training distribution and evaluate its performance on many unseen holdout tasks. (b) To construct a training distribution, we sample an initial state from a predefined distribution, and run a goal setting policy (Alice) to generate a goal. In one episode, Alice is asked to generate 5 goals and Bob solves them in sequence until it fails.

et al., 2018; Nair et al., 2018b; Racaniere et al., 2020) and collecting teleoperated robot trajectories to identify goals (Lynch et al., 2020; Gupta et al., 2020). In this paper, we extend an alternative approach called asymmetric self-play (Sukhbaatar et al., 2018b;a) for automated goal generation. Asymmetric self-play trains two RL agents named Alice and Bob. Alice learns to propose goals that Bob is likely to fail at, and Bob, a goal-conditioned policy, learns to solve the proposed goals. Alice proposes a goal by manipulating objects and Bob has to solve the goal starting from the same initial state as Alice's. By embodying these two agents into the same robotic hardware, this setup ensures that all proposed goals are provided with at least one solution: Alice's trajectory.

There are two main reasons why we consider asymmetric self-play to be a promising goal generation and learning method. First, any proposed goal is *achievable*, meaning that there exists at least one solution trajectory that Bob can follow to achieve the goal. Because of this property, we can exploit Alice's trajectory to provide additional learning signal to Bob via behavioral cloning. This additional learning signal alleviates the overhead of heuristically designing a curriculum or reward shaping for learning. Second, this approach does not require labor intensive data collection.

In this paper, we show that asymmetric self-play can be used to train a goal-conditioned policy for complex object manipulation tasks, and the learned policy can zero-shot generalize to many manually designed holdout tasks, which consist of either previously unseen goals, previously unseen objects, or both. To the best of our knowledge, this is the first work that presents zero-shot generalization to many previously unseen tasks by training purely with asymmetric self-play.[1]

2 Problem Formulation

Our training environment for robotic manipulation consists of a robot arm with a gripper attached and a wide range of objects placed on a table surface (Figure 1a,1b). The goal-conditioned policy learns to control the robot to rearrange randomly placed objects (the initial state) into a specified goal configuration (Figure 1c). We aim to train a policy on a single training distribution and to evaluate its performance over a suite of holdout tasks which are independently designed and not explicitly present during training (Figure 2a). In this work, we construct the training distribution via *asymmetric self-play* (Figure 2b) to achieve generalization to many unseen holdout tasks (Figure 1c).

Mathematical formulation Formally, we model the interaction between an environment and a goal-conditioned policy as a goal-augmented Markov decision process $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \mathcal{G} \rangle$, where $\mathcal{S}$ is the state space, $\mathcal{A}$ is the action space, $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \mapsto \mathbb{R}$ denotes the transition probability, $\mathcal{G} \subseteq \mathcal{S}$ specifies the goal space and $\mathcal{R} : \mathcal{S} \times \mathcal{G} \mapsto \mathbb{R}$ is a goal-specific reward function. A goal-augmented trajectory sequence is $\{(s_0, g, a_0), \ldots, (s_t, g, a_t)\}$, where the goal is provided to the policy as part of the observation at every step. We say a goal is achieved if s_t is sufficiently close to g (Appendix A.2). With a slightly overloaded notation, we define the *goal distribution* $\mathcal{G}(g|s_0)$ as the probability of a goal state $g \in \mathcal{G}$ conditioned on an initial state $s_0 \in \mathcal{S}$.

[1]Asymmetric self-play is proposed in Sukhbaatar et al. (2018b;a), but to supplement training while the majority of training is conducted on target tasks. Zero-shot generalization to unseen tasks was not evaluated.

Training goal distribution A naive design of the goal distribution $\mathcal{G}(g|s_0)$ is to randomly place objects uniformly on the table, but it is unlikely to generate interesting goals, such as an object picked up and held above the table surface by a robot gripper. Another possible approach, collecting tasks and goals manually, is expensive and hard to scale. We instead sidestep these issues and automatically generate goals via training based on asymmetric self-play (Sukhbaatar et al., 2018b;a). Asymmetric self-play involves using a policy named Alice $\pi_A(a|s)$ to set goals and a goal-conditioned policy Bob $\pi_B(a|s, g)$ to solve goals proposed by Alice, as illustrated in Figure 2b. We run π_A to generate a trajectory $\{(s_0, a_0), \ldots, (s_T, a_T)\}$ and the last state is labelled as a goal g for π_B to solve. The goal distribution $\mathcal{G}(s_T = g|s_0)$ is fully determined by π_A and we train Bob only on this goal distribution. We therefore say *zero-shot generalization* when Bob generalizes to a holdout task which is not explicitly encoded into the training distribution.

Evaluation on holdout tasks To assess zero-shot generalization of $\pi_B(a|s, g)$ from our training setup, we hand-designed a suite of holdout tasks with goals that are never directly incorporated into the training distribution. Some holdout tasks also feature previously unseen objects. The holdout tasks are designed to either test whether a specific skill has been learned, such as the ability to pick up objects (Figure 3), or represent a semantically interesting task, such as setting a table (Figure 1c). Appendix B.6 describes the list of holdout tasks that we use in our experiments.

3 Asymmetric Self-play

To train Alice policy $\pi_A(a|s)$ and Bob policy $\pi_B(a|s, g)$, we run the following multi-goal game within one episode, as illustrated in Figure 2b:

1. An initial state s_0 is sampled from an initial state distribution. Alice and Bob are instantiated into their own copies of the environment. Alice and Bob alternate turns as follows.
2. **Alice's turn.** Alice interacts with its environment for a fixed number of T steps and may rearrange the objects. The state at the end of Alice's turn s_T will be used as a goal g for Bob. If the proposed goal is invalid (e.g. if Alice has not moved any objects, or if an object has fallen off the table), the episode terminates.
3. **Bob's turn.** Bob receives reward if it successfully achieves the goal g in its environment. Bob's turn ends when it succeeds at achieving the goal or reaches a timeout. If Bob's turn ends in a failure, its remaining turns are skipped and treated as failures, while we let Alice to keep generating goals.
4. Alice receives reward if Bob fails to solve the goal that Alice proposed. Steps 2–3 are repeated until 5 goals are set by Alice or Alice proposes an invalid goal, and then the episode terminates.

The competition created by this game encourages Alice to propose goals that are increasingly challenging to Bob, while Bob is forced to solve increasingly complex goals. The multi-goal setup was chosen to allow Bob to take advantage of environmental information discovered earlier in the episode to solve its remaining goals, which OpenAI et al. (2019a) found to be important for transfer to physical systems. Note however that in this work we focus on solving goals in simulation only. To improve stability and avoid forgetting, we have Alice and Bob play against past versions of their respective opponent in 20% of games. More details about the game structure are in Appendix A.

3.1 Reward Structure

For Bob, we assign *sparse* goal-conditioned rewards. Whenever an object is placed at its goal position and orientation up to a small error (the *success threshold*), Bob receives 1 reward immediately. But if this object is moved away from the goal state that it has arrived at in past steps, Bob obtains -1 reward such that the sum of per-object reward is at most 1 during a given turn. When all of the objects are in their goal state, Bob receives 5 additional reward and its turn is over.

For Alice, we assign a reward after Bob has attempted to solve the goal: 5 reward if Bob failed at solving the goal, and 0 if Bob succeeded. We shape Alice's reward slightly by adding 1 reward if it has set a valid goal, defined to be when no object has fallen off the table and any object has been moved more than the success threshold. An additional penalty of -3 reward is introduced when

Alice sets a goal with objects outside of the placement area, defined to be a fixed 3D volume within the view of the robot's camera. More details are discussed in Appendix A.2.

3.2 ALICE BEHAVIORAL CLONING (ABC)

One of the main benefits of using asymmetric self-play is that the generated goals come with at least one solution to achieve it: *Alice's trajectory*. Similarly to Sukhbaatar et al. (2018a), we exploit this property by training Bob with Behavioral Cloning (BC) from Alice's trajectory, in addition to the reinforcement learning (RL) objective. We call this learning mechanism *Alice Behavioral Cloning* (ABC). We propose several improvements over the original formulation in Sukhbaatar et al. (2018a).

Demonstration trajectory filtering Compared to BC from expert demonstrations, using Alice's trajectory needs extra care. Alice's trajectory is likely to be suboptimal for solving the goal, as Alice might arrive at the final state merely by accident. Therefore, we only consider trajectories with goals that Bob failed to solve as demonstrations, to avoid distracting Bob with suboptimal examples. Whenever Bob fails, we relabel Alice's trajectory $\{(s_0, a_0), \ldots, (s_T, a_T)\}$ to be a goal-augmented version $\{(s_0, s_T, a_0), \ldots, (s_T, s_T, a_T)\}$ as a demonstration for BC, where s_T is the goal.

PPO-style BC loss clipping The objective for training Bob is $\mathcal{L} = \mathcal{L}_{\text{RL}} + \beta \mathcal{L}_{\text{abc}}$, where $\mathcal{L}_{\text{RL}}$ is an RL objective and $\mathcal{L}_{\text{abc}}$ is the ABC loss. β is a hyperparameter controlling the relative importance of the BC loss. We set $\beta = 0.5$ throughout the whole experiment. A naive BC loss is to minimize the negative log-likelihood of demonstrated actions, $-\mathbb{E}_{(s_t, g_t, a_t) \in \mathcal{D}_{\text{BC}}} \big[\log \pi_B(a_t | s_t, g_t; \theta) \big]$ where $\mathcal{D}_{\text{BC}}$ is a mini-batch of demonstration data and π_B is parameterized by θ. We found that overly-aggressive policy changes triggered by BC sometimes led to learning instabilities. To prevent the policy from changing too drastically, we introduce PPO-style loss clipping (Schulman et al., 2017) on the BC loss by setting the advantage $\hat{A} = 1$ in the clipped surrogate objective:

$$\mathcal{L}_{\text{abc}} = -\mathbb{E}_{(s_t, g_t, a_t) \in \mathcal{D}_{\text{BC}}} \left[\text{clip}\Big(\frac{\pi_B(a_t | s_t, g_t; \theta)}{\pi_B(a_t | s_t, g_t; \theta_{\text{old}})}, 1 - \epsilon, 1 + \epsilon \Big) \right]$$

where $\pi_B(a_t | s_t, g_t; \theta)$ is Bob's likelihood on a demonstration based on the parameters that we are optimizing, and $\pi_B(a_t | s_t, g_t; \theta_{\text{old}})$ is the likelihood based on Bob's behavior policy (at the time of demonstration collection) evaluated on a demonstration. This behavior policy is identical to the policy that we use to collect RL trajectories. By setting $\hat{A} = 1$, this objective optimizes the naive BC loss, but clips the loss whenever $\frac{\pi_B(a_t | s_t, g_t; \theta)}{\pi_B(a_t | s_t, g_t; \theta_{\text{old}})}$ is bigger than $1 + \epsilon$, to prevent the policy from changing too much. ϵ is a clipping threshold and we use $\epsilon = 0.2$ in all the experiments.

4 RELATED WORK

Training distribution for RL In the context of multi-task RL (Beattie et al., 2016; Hausman et al., 2018; Yu et al., 2020), multi-goal RL (Kaelbling, 1993; Andrychowicz et al., 2017), and meta RL (Wang et al., 2016; Duan et al., 2016), previous works manually designed a distribution of tasks or goals to see better generalization of a policy to a new task or goal. Domain randomization (Sadeghi & Levine, 2017; Tobin et al., 2017; OpenAI et al., 2020) manually defines a distribution of simulated environments, but in service of generalizing to the same task in the real world.

There are approaches to grow the training distribution automatically (Srivastava et al., 2013). Self-play (Tesauro, 1995; Silver et al., 2016; 2017; Bansal et al., 2018; OpenAI et al., 2019b; Vinyals et al., 2019) constructs an ever-growing training distribution where multiple agents learn by competing with each other, so that the resulting agent shows strong performance on a single game. OpenAI et al. (2019a) automatically grew a distribution of domain randomization parameters to accomplish better generalization in the task of solving a Rubik's cube on the physical robot. Wang et al. (2019; 2020) studied an automated way to keep discovering challenging 2D terrains and locomotion policies that can solve them in a 2D bipedal walking environment.

We employ asymmetric self-play to construct a training distribution for learning a goal-conditioned policy and to achieve generalization to unseen tasks. Florensa et al. (2018); Nair et al. (2018b); Racaniere et al. (2020) had the same motivation as ours, but trained a generative model instead of

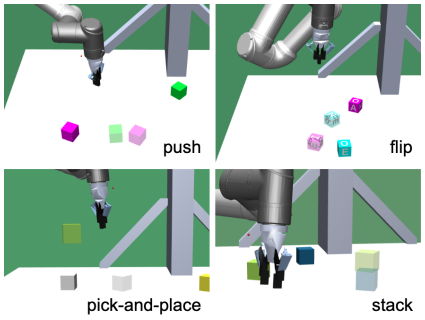

Figure 3: Holdout tasks in the environment using 1 or 2 blocks. The transparent blocks denote the desired goal state, while opaque blocks are the current state. (a) `push`: The blocks must be moved to their goal locations and orientations. There is no differentiation between the six block faces. (b) `flip`: Each side of the block is labelled with a unique letter. The blocks must be moved to make every face correctly positioned as what the goal specifies. (c) `pick-and-place`: One goal block is in the air. (d) `stack`: Two blocks must be stacked in the right order at the right location.

a goal setting policy. Thus, the difficulties of training a generative model were inherited by these methods: difficulty of modeling a high dimensional space and generation of unrealistic samples. Lynch et al. (2020); Gupta et al. (2020) used teleoperation to collect arbitrary robot trajectories, and defined a goal distribution from the states in the collected trajectories. This approach likely requires a large number of robot trajectories for each environment configuration (e.g. various types of objects on a table), and randomization of objects was not studied in this context.

Asymmetric self-play Asymmetric self-play was proposed by Sukhbaatar et al. (2018b;a) as a way to supplement RL training. Sukhbaatar et al. (2018b) mixed asymmetric self-play training with standard RL training on the target task and measure the performance on the target task. Sukhbaatar et al. (2018a) used asymmetric self-play to pre-train a hierarchical policy, and evaluated the policy after finetuning on a target task. As far as we know, no previous work has trained a goal-conditioned policy *purely* based on asymmetric self-play and evaluated generalization to unseen holdout tasks.

Curriculum learning Many previous works showed the difficulty of RL and proposed an automated curriculum (Andrychowicz et al., 2017; Florensa et al., 2017; Salimans & Chen, 2018; Matiisen et al., 2019; Zhang et al., 2020) or auxiliary exploration objectives (Oudeyer et al., 2007; Baranes & Oudeyer, 2013; Pathak et al., 2017; Burda et al., 2019; Ecoffet et al., 2019; 2020) to learn *predefined tasks*. Our work shares a similarity in that asymmetric self-play alleviates the difficulty of learning a goal-conditioned policy via an intrinsic curriculum and imitation from the goal setter's trajectory, but our work does not assume any predefined task or goal distribution.

Robot learning for object manipulation. It has been reported that training a policy for multi-object manipulation is very challenging with *sparse* rewards (Riedmiller et al., 2018; Vecerik et al., 2018). One example is block stacking, which has been studied for a long time in robotics as it involves complex contact reasoning and long horizon motion planning (Deisenroth et al., 2011). Learning block stacking often requires a hand-designed curriculum (Li et al., 2019), meticulous reward shaping (Popov et al., 2017), fine-tuning (Rusu et al., 2017), or human demonstrations (Nair et al., 2018a; Duan et al., 2017). In this work, we use block stacking as one of the holdout tasks to test zero-shot generalization, but without training on it.

5 Experiments

In this section, we first show that asymmetric self-play generates an effective training curriculum that enables generalization to unseen hold-out tasks. Then, we scale up the experiment to train in an environment containing multiple random complex objects and evaluate it with a set of holdout tasks containing unseen objects and unseen goal configurations. Finally, we demonstrate how critical ABC is for Bob to make progress in a set of ablation studies.

5.1 Experimental setup

We implement the training environment described in Sec. 2 with randomly placed ShapeNet objects (Chang et al., 2015) as an initial state distribution. In addition, we set up another simpler environment using one or two blocks of fixed size, used for small-scale comparisons and ablation studies. Figure 3 visualizes four holdout tasks for this environment. Each task is designed to evaluate whether the robot has acquired certain manipulation skills: pushing, flipping, picking up and

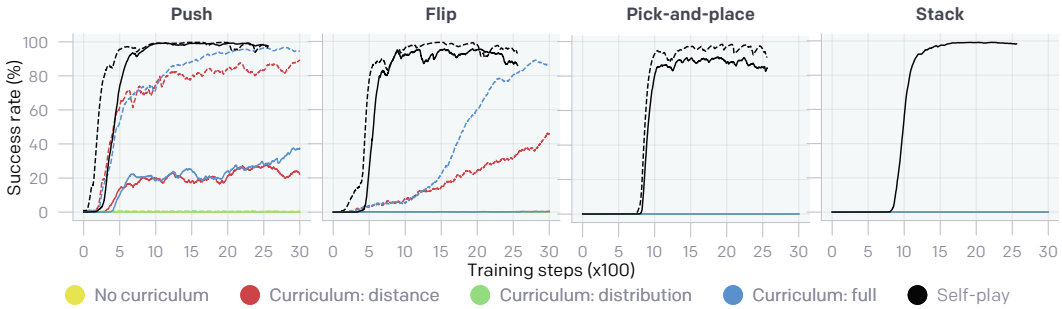

Figure 4: Generalization to unseen holdout tasks for blocks. Baselines are trained over a mixture of all holdout tasks. The solid lines represent 2-blocks, while the dashed lines are for 1-block.

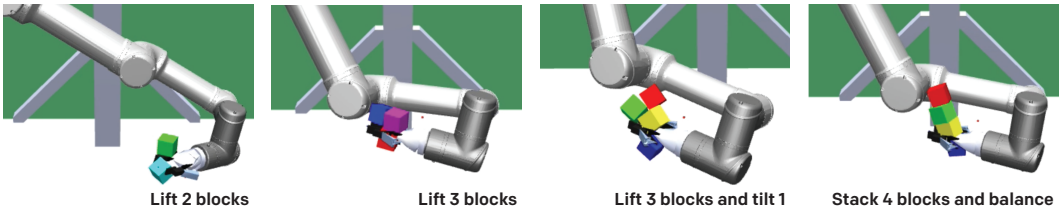

Figure 5: Goals discovered by asymmetric self-play. Alice discovers many goals that are not covered by our manually designed holdout tasks on blocks.

stacking blocks. Experiments in Sec. 5.2, 5.3 and 5.5 focus on blocks and experimental results based on ShapeNet objects are present on Sec. 5.4. More details on our training setups are in Appendix B.

We implement Alice and Bob as two independent policies of the same network architecture with memory (Appendix B.4), except that Alice has no observation on goal state. The policies take state observations ("state policy") for experiments with blocks (Sec. 5.2, 5.3, and 5.5), and take both vision and state observations ("hybrid policy") for experiments with ShapeNet objects (Sec. 5.4). Both policies are trained with Proximal Policy Optimization (PPO) (Schulman et al., 2017).

5.2 Generalization to unseen goals without manual curricula

One way to train a single policy to acquire all the skills in Figure 3 is to train a goal-conditioned policy directly over a mixture of these tasks. However, training directly over these tasks without a curriculum turns out to be very challenging, as the policy completely fails to make any progress.[2] In contrast, Bob is able to solve all these holdout tasks quickly when learning via asymmetric self-play, without explicitly encoding any prior knowledge of the holdout tasks into the training distribution.

To gauge the effect of an intrinsic curriculum introduced by self-play, we carefully designed a set of non-self-play baselines using explicit curricula controlled by Automatic Domain Randomization (OpenAI et al., 2019a). All baselines are trained over a mixture of block holdout tasks as the goal distribution. We measure the effectiveness of a training setup by tracking the success rate for each holdout task, as shown in Figure 4. The `no curriculum` baseline fails drastically. The `curriculum:distance` baseline expands the distance between the initial and goal states gradually as training progresses, but only learns to push and flip a single block. The `curriculum:distribution` baseline, which slowly increases the proportion of pick-and-place and stacking goals in the training distribution, fails to acquire any skill. The `curriculum:full` baseline incorporates all hand-designed curricula yet still cannot learn how to pick up or stack blocks. We have spent a decent amount of time iterating and improving these baselines but found it especially difficult to develop a scheme good enough to compete with asymmetric self-play. See Appendix C.1 for more details of our baselines.

5.3 Discovery of novel goals and solutions

Asymmetric self-play discovers novel goals and solutions that are not covered by our holdout tasks. As illustrated in Figure 5, Alice can lift multiple blocks at the same time, build a tower and then

[2]The tasks was easier when we ignored object rotation as part of the goal, and used a smaller table.

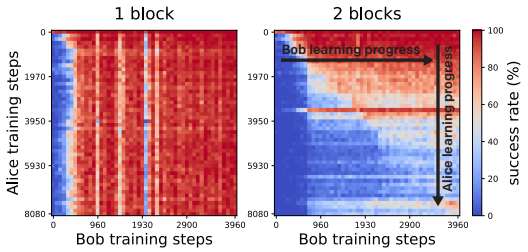

Figure 6: The empirical payoff matrix between Alice and Bob. Average success rate over multiple self-play episode is visualized. Alice with more training steps generates more challenging goals that Bob cannot solve yet. Bob with more training steps can achieve more goals against the same Alice.

keep it balanced using an arm joint. Although it is a tricky strategy for Bob to learn on its own, with ABC, Bob eventually acquires the skills for solving such complex tasks proposed by Alice. Videos are available at https://robotics-self-play.github.io.

Figure 6 summarizes Alice and Bob's learning progress against each other. For every pair of Alice and Bob, we ran multiple self-play episodes and measured the success rate. We observe an interesting trend with 2 blocks. As training proceeds, Alice tends to generate more challenging goals, where Bob shows lower success rate. With past sampling, Bob continues to make progress against versions of Alices from earlier optimization steps. This visualization suggests a desired dynamic of asymmetric self-play that could potentially lead to unbounded complexity: Alice continuously generates goals to challenge Bob, and Bob keeps making progress on learning to solve new goals.

5.4 Generalization to unseen objects and goals

The experiments above show strong evidence that efficient curricula and novel goals can autonomously emerge in asymmetric self-play. To further challenge our approach, we scale it up to work with many more complex objects using more computational resources for training. We train a hybrid policy in an environment containing up to 10 random ShapeNet (Chang et al., 2015) objects. During training, we randomize the number of objects and the object sizes via Automatic Domain Randomization (OpenAI et al., 2019a). The hybrid policy uses vision observations to extract information about object geometry and size. We evaluate the Bob policy on a more diverse set of manipulation tasks, including semantically interesting ones. Many tasks contain unseen objects and complex goals, as illustrated in Figure 7.

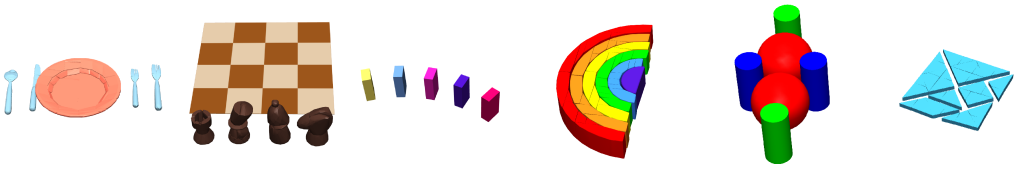

(a) Table setting (b) Mini chess (c) Domino (d) Rainbow (e) Ball-capture (f) Tangram

Figure 7: Example holdout tasks involving unseen objects and complex goal states. The goal states are illustrated here, and the initial states have randomly placed objects.

The learned Bob policy achieves decent zero-shot generalization performance for many tasks. Success rates are reported in Figure 8. Several tasks are still challenging. For example, `ball-capture` requires delicate handling of rolling objects and lifting skills. The `rainbow` tasks call for an understanding of concave shapes. Understanding the ordering of placement actions is crucial for stacking more than 3 blocks in the desired order. The Bob policy learns such an ordering to some degree, but fails to fully generalize to an arbitrary number of stacked blocks.

5.5 Ablation studies

We present a series of ablation studies designed for measuring the importance of each component in our asymmetric self-play framework, including Alice behavioral cloning (ABC), BC loss clipping, demonstration filtering, and the multi-goal game setup. We disable a single ingredient in each ablation run and compare with the complete self-play baseline in Figure 9.

The `no ABC` baseline shows that Bob *completely fails* to solve any holdout task without ABC, indicating that ABC is a critical mechanism in asymmetric self-play. The `no BC loss clipping` baseline

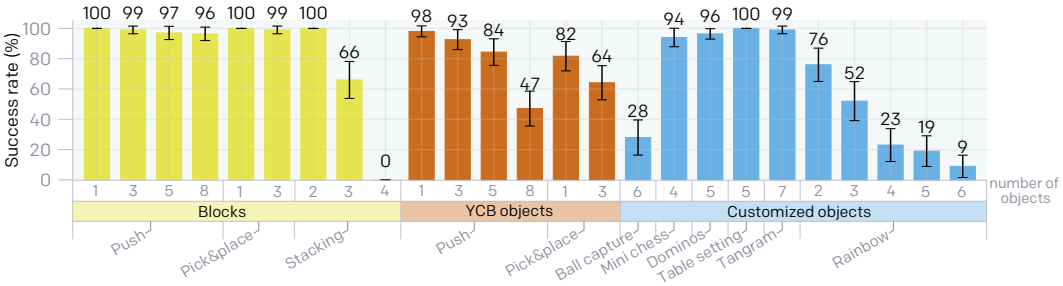

Figure 8: Success rates of a single goal-conditioned policy solving a variety of holdout tasks, averaged over 100 trials. The error bars indicate the 99% confidence intervals. Yellow, orange and blue bars correspond to success rates of manipulation tasks with blocks, YCB[3] objects and other uniquely built objects, respectively. Videos are available at https://robotics-self-play.github.io.

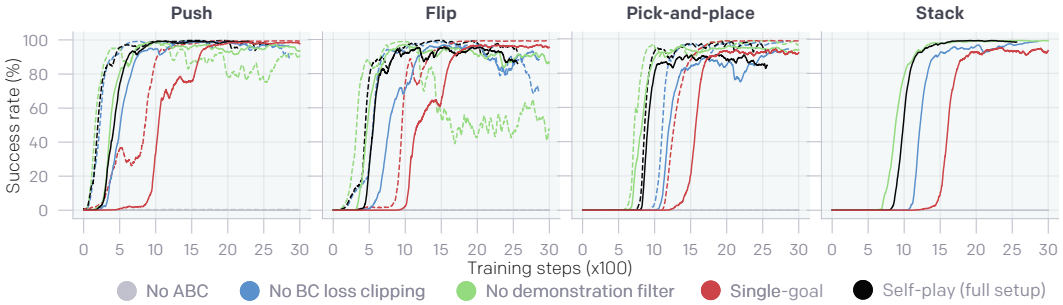

Figure 9: The ablation studies compare four ablation runs each with one component disabled with the full baseline. Solid lines are for 2-blocks, dashed lines are for 1-block.

shows slightly slower learning on `pick-and-place` and `stack`, as well as some instabilities in the middle of training. Clipping in the BC loss is expected to help alleviate this instability by controlling the rate of policy change per optimizer iteration. The `no demonstration filter` baseline shows noticeable instability on `flip`, suggesting the importance of excluding suboptimal demonstrations from behavioral cloning. Finally, the `single-goal` baseline uses a single goal instead of 5 goals per episode during training. The evaluation tasks are also updated to require a single success per episode. Generalization of this baseline to holdout tasks turns out to be much slower and less stable. It signifies some advantages of using multiple goals per episode, perhaps due to the policy memory internalizing environmental information during multiple trials of goal solving.

The results of the ablation studies suggest that ABC with proper configuration and multi-goal gameplay are critical components of asymmetric self-play, alleviating the importance of manual curricula and facilitating efficient learning.

6 Conclusion

We studied asymmetric self-play as a framework for defining a single training distribution to learn many arbitrary object manipulation tasks. Even without any prior knowledge about the target tasks, asymmetric self-play is able to train a strong goal-conditioned policy that can generalize to many unseen holdout tasks. We found that asymmetric self-play not only generates a wide range of interesting goals but also alleviates the necessity of designing manual curricula for learning such goals. We provided evidence that using the goal setting trajectory as a demonstration for training a goal solving policy is essential to enable efficient learning. We further scaled up our approach to work with various complex objects using more computation, and achieved zero-shot generalization to a collection of challenging manipulation tasks involving unseen objects and unseen goals.

[3] https://www.ycbbenchmarks.com/object-models/

REFERENCES

Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, Pieter Abbeel, and Wojciech Zaremba. Hindsight experience replay. In *Advances in neural information processing systems*, 2017.

Trapit Bansal, Jakub Pachocki, Szymon Sidor, Ilya Sutskever, and Igor Mordatch. Emergent complexity via multi-agent competition. In *International Conference on Learning Representations*, 2018.

Adrien Baranes and Pierre-Yves Oudeyer. Active learning of inverse models with intrinsically motivated goal exploration in robots. *Robotics and Autonomous Systems*, 61(1):49–73, 2013.

Charles Beattie, Joel Z Leibo, Denis Teplyashin, Tom Ward, Marcus Wainwright, Heinrich Küttler, Andrew Lefrancq, Simon Green, Víctor Valdés, Amir Sadik, et al. Deepmind lab. *arXiv preprint arXiv:1612.03801*, 2016.

Yuri Burda, Harrison Edwards, Amos Storkey, and Oleg Klimov. Exploration by random network distillation. In *International Conference on Learning Representations*, 2019.

Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.

Marc Peter Deisenroth, Carl Edward Rasmussen, and Dieter Fox. Learning to control a low-cost manipulator using data-efficient reinforcement learning. *Robotics: Science and Systems VII*, pp. 57–64, 2011.

Yan Duan, John Schulman, Xi Chen, Peter L Bartlett, Ilya Sutskever, and Pieter Abbeel. RL^2: Fast reinforcement learning via slow reinforcement learning. *arXiv preprint arXiv:1611.02779*, 2016.

Yan Duan, Marcin Andrychowicz, Bradly Stadie, OpenAI Jonathan Ho, Jonas Schneider, Ilya Sutskever, Pieter Abbeel, and Wojciech Zaremba. One-shot imitation learning. In *Advances in neural information processing systems*, 2017.

Adrien Ecoffet, Joost Huizinga, Joel Lehman, Kenneth O Stanley, and Jeff Clune. Go-explore: a new approach for hard-exploration problems. *arXiv preprint arXiv:1901.10995*, 2019.

Adrien Ecoffet, Joost Huizinga, Joel Lehman, Kenneth O Stanley, and Jeff Clune. First return then explore. *arXiv preprint arXiv:2004.12919*, 2020.

Lasse Espeholt, Hubert Soyer, Remi Munos, Karen Simonyan, Vlad Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, et al. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures. In *International Conference on Machine Learning*, 2018.

Carlos Florensa, David Held, Markus Wulfmeier, Michael Zhang, and Pieter Abbeel. Reverse curriculum generation for reinforcement learning. In *Conference on Robot Learning*, 2017.

Carlos Florensa, David Held, Xinyang Geng, and Pieter Abbeel. Automatic goal generation for reinforcement learning agents. In *International conference on machine learning*, 2018.

Shixiang Gu, Ethan Holly, Timothy Lillicrap, and Sergey Levine. Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In *IEEE international conference on robotics and automation (ICRA)*, 2017.

Abhishek Gupta, Vikash Kumar, Corey Lynch, Sergey Levine, and Karol Hausman. Relay policy learning: Solving long-horizon tasks via imitation and reinforcement learning. In *Conference on Robot Learning*, 2020.

Karol Hausman, Jost Tobias Springenberg, Ziyu Wang, Nicolas Heess, and Martin Riedmiller. Learning an embedding space for transferable robot skills. In *International Conference on Learning Representations*, 2018.

Jemin Hwangbo, Joonho Lee, Alexey Dosovitskiy, Dario Bellicoso, Vassilios Tsounis, Vladlen Koltun, and Marco Hutter. Learning agile and dynamic motor skills for legged robots. *Science Robotics*, 4(26), 2019.

Leslie Pack Kaelbling. Learning to achieve goals. In *Proceedings of the 13th International Joint Conference on Artificial Intelligence*, 1993.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

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.

Richard Li, Allan Jabri, Trevor Darrell, and Pulkit Agrawal. Towards practical multi-object manipulation using relational reinforcement learning. *arXiv preprint arXiv:1912.11032*, 2019.

Corey Lynch, Mohi Khansari, Ted Xiao, Vikash Kumar, Jonathan Tompson, Sergey Levine, and Pierre Sermanet. Learning latent plans from play. In *Conference on Robot Learning*, pp. 1113–1132, 2020.

Tambet Matiisen, Avital Oliver, Taco Cohen, and John Schulman. Teacher-student curriculum learning. *IEEE transactions on neural networks and learning systems*, 2019.

Ashvin Nair, Bob McGrew, Marcin Andrychowicz, Wojciech Zaremba, and Pieter Abbeel. Overcoming exploration in reinforcement learning with demonstrations. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2018a.

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*, 2018b.

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. *arXiv preprint arXiv:1910.07113*, 2019a.

OpenAI, Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemysław Dębiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, Rafal Józefowicz, Scott Gray, Catherine Olsson, Jakub Pachocki, Michael Petrov, Henrique Pondé de Oliveira Pinto, Jonathan Raiman, Tim Salimans, Jeremy Schlatter, Jonas Schneider, Szymon Sidor, Ilya Sutskever, Jie Tang, Filip Wolski, and Susan Zhang. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*, 2019b.

OpenAI, Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Jozefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, Jonas Schneider, Szymon Sidor, Josh Tobin, Peter Welinder, Lilian Weng, and Wojciech Zaremba. Learning dexterous in-hand manipulation. *The International Journal of Robotics Research*, 39(1):3–20, 2020.

Pierre-Yves Oudeyer, Frdric Kaplan, and Verena V Hafner. Intrinsic motivation systems for autonomous mental development. *IEEE transactions on evolutionary computation*, 11(2):265–286, 2007.

Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 2017.

Ivaylo Popov, Nicolas Heess, Timothy Lillicrap, Roland Hafner, Gabriel Barth-Maron, Matej Vecerik, Thomas Lampe, Yuval Tassa, Tom Erez, and Martin Riedmiller. Data-efficient deep reinforcement learning for dexterous manipulation. *arXiv preprint arXiv:1704.03073*, 2017.

Sebastien Racaniere, Andrew Lampinen, Adam Santoro, David Reichert, Vlad Firoiu, and Timothy Lillicrap. Automated curriculum generation through setter-solver interactions. In *International Conference on Learning Representations*, 2020.

Martin Riedmiller, Roland Hafner, Thomas Lampe, Michael Neunert, Jonas Degrave, Tom Van de Wiele, Volodymyr Mnih, Nicolas Heess, and Jost Tobias Springenberg. Learning by playing - solving sparse reward tasks from scratch, 2018.

Andrei A Rusu, Matej Večerík, Thomas Rothörl, Nicolas Heess, Razvan Pascanu, and Raia Hadsell. Sim-to-real robot learning from pixels with progressive nets. In *Conference on Robot Learning*, 2017.

Fereshteh Sadeghi and Sergey Levine. Cad2rl: Real single-image flight without a single real image. In *Proceedings of Robotics: Science and Systems*, 2017.

Tim Salimans and Richard Chen. Learning montezuma's revenge from a single demonstration. *arXiv preprint arXiv:1812.03381*, 2018.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.

David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, et al. Mastering chess and shogi by self-play with a general reinforcement learning algorithm. *arXiv preprint arXiv:1712.01815*, 2017.

Rupesh Kumar Srivastava, Bas R Steunebrink, and Jürgen Schmidhuber. First experiments with powerplay. *Neural Networks*, 41:130–136, 2013.

Sainbayar Sukhbaatar, Emily Denton, Arthur Szlam, and Rob Fergus. Learning goal embeddings via self-play for hierarchical reinforcement learning. *arXiv preprint arXiv:1811.09083*, 2018a.

Sainbayar Sukhbaatar, Zeming Lin, Ilya Kostrikov, Gabriel Synnaeve, Arthur Szlam, and Rob Fergus. Intrinsic motivation and automatic curricula via asymmetric self-play. In *International Conference on Learning Representations*, 2018b.

Gerald Tesauro. Temporal difference learning and td-gammon. *Communications of the ACM*, 38(3): 58–68, 1995.

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. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017.

Emanuel Todorov, Tom Erez, and Yuval Tassa. MuJoCo: A physics engine for model-based control. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, 2012.

Mel Vecerik, Todd Hester, Jonathan Scholz, Fumin Wang, Olivier Pietquin, Bilal Piot, Nicolas Heess, Thomas Rothörl, Thomas Lampe, and Martin Riedmiller. Leveraging demonstrations for deep reinforcement learning on robotics problems with sparse rewards, 2018.

Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.

Jane X Wang, Zeb Kurth-Nelson, Dhruva Tirumala, Hubert Soyer, Joel Z Leibo, Remi Munos, Charles Blundell, Dharshan Kumaran, and Matt Botvinick. Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763*, 2016.

Rui Wang, Joel Lehman, Jeff Clune, and Kenneth O Stanley. Paired open-ended trailblazer (poet): Endlessly generating increasingly complex and diverse learning environments and their solutions. *arXiv preprint arXiv:1901.01753*, 2019.

Rui Wang, Joel Lehman, Aditya Rawal, Jiale Zhi, Yulun Li, Jeff Clune, and Kenneth O Stanley. Enhanced poet: Open-ended reinforcement learning through unbounded invention of learning challenges and their solutions. *arXiv preprint arXiv:2003.08536*, 2020.

Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on Robot Learning*, 2020.

Yunzhi Zhang, Pieter Abbeel, and Lerrel Pinto. Automatic curriculum learning through value disagreement. *arXiv preprint arXiv:2006.09641*, 2020.