

# Achieving Goals using Reward Shaping and Curriculum Learning

Mihai Anca<sup>1</sup>, Jonathan D. Thomas<sup>1</sup>, Dabal Pedamonti<sup>1</sup>, Mark Hansen<sup>2</sup>, and Matthew Studley<sup>2</sup>

<sup>1</sup> University of Bristol, Bristol, United Kingdom  
`mihai.anca@bristol.ac.uk`

<sup>2</sup> University of the West of England, Bristol, United Kingdom

**Abstract.** Real-time control for robotics is a popular research area in the reinforcement learning community. Through the use of techniques such as reward shaping, researchers have managed to train online agents across a multitude of domains. Despite these advances, solving goal-oriented tasks still requires complex architectural changes or hard constraints to be placed on the problem. In this article, we solve the problem of stacking multiple cubes by combining curriculum learning, reward shaping, and a high number of efficiently parallelized environments. We introduce two curriculum learning settings that allow us to separate the complex task into sequential sub-goals, hence enabling the learning of a problem that may otherwise be too difficult. We focus on discussing the challenges encountered while implementing them in a goal-conditioned environment. Finally, we extend the best configuration identified on a higher complexity environment with differently shaped objects.

**Keywords:** reinforcement learning, curriculum learning, reward shaping, robotics

## 1 Introduction

Long-horizon manipulation tasks, such as object stacking, require planning over an extended period of time while taking into account various mandatory checkpoints. Creating agents that can solve this type of task using complex reward signals is a key challenge in deep reinforcement learning (deep RL). Recent advances have shown that RL applied to robotic arms can perform object stacking [11, 10, 14, 16], however, the problem is usually constrained, or the solution requires complex changes to the model’s architecture. Moreover, standard RL approaches do not scale well when the number of task objects is increased.

We believe that a curriculum learning (CL) framework can help alleviate this issue by decomposing the task into sub-goals of incremental difficulty. This has been shown by Narvekar et al. [20] to be effective in applications such as robotics and games.

We leveraged the use of the simulation platform Isaac Gym [18], which allows thousands of environments to be asynchronously controlled by an RL agent. We

found that there is a gap in the on-policy RL literature when dealing with heavily parallelized object-centric goal-oriented tasks, such as stacking. Our work, to the best of our knowledge, provides the first results in this territory.

Given that we used an online on-policy algorithm, the methodology and findings described in this paper can be transferred to other sample-efficient algorithms that train agents directly in the real world.

In summary, the paper makes the following contributions: 1) We present various CL setups and couple them with reward shaping, enabling the stacking of three objects using standard proximal policy optimization (PPO) [26]. To the best of our knowledge, no prior work successfully achieved this task using on-policy vanilla RL algorithms (i.e. PPO). 2) We analyse the elements that hinder learning and discuss the key components that are critical for task completion. 3) We show that in using an asynchronous setup, such as Isaac Gym [18], to learn goal-oriented tasks, the learning horizon becomes an important hyperparameter. We then consider the effects of a curriculum for adjusting this value. 4) Finally, we provide the *codebase* to the community to act as a baseline for future work in multiple cube-stacking tasks.

**A surprising finding:** Choosing the optimal learning horizon value for the whole task of stacking three cubes results in a better convergence time when compared to selecting the optimal value per curriculum stage. Allowing 50 timesteps for placing each object, the agent trains faster and achieves a higher accuracy if trained from the start with a horizon length of 150 (assuming three objects in the scene). This finding is discussed in Section 5.

## 2 Background and Related Work

The “Pick and Place” task was popularised as part of a reinforcement learning (RL) baseline with the release of OpenAI robotics gym pack [4] in 2016. Both online and offline RL methods have been shown to work on the dense reward version of this environment. For example, in Jeong et al. [11], the authors used a model-based vision approach and self-supervision to pick and place one cube on top of another.

However, the literature to date offers no simple solution to the problem of stacking at least three objects, which represents a key example of a long-horizon task. In Hundt et al. [10], the authors discretised the action space and managed to learn an RL policy that successfully stacks multiple objects. This method relies on prior knowledge and is only applicable to simple-to-grasp objects, such as cubes. The recent work of Lee et al. [14] showed a potential solution when stacking differently shaped objects, however, the problem solved is heavily constrained and requires multiple training sessions.

This problem becomes increasingly difficult when posed in a sparse reward environment. Li et al. [16] overcame it by using graph neural networks (GNN) coupled with attention layers, successfully stacking up to six cubes in simulation. However, even though the number of trainable parameters is similar to a fully connected network, the new architecture requires considerably more VRAM

when computing the message-parsing step of the GNN. Additionally, it failed to explain the underlying issue of why vanilla RL approaches are incapable of learning this task.

Another popular example of stacking multiple cubes using sparse rewards is that presented in Nair et al. [19]. The authors made clever use of the critic network in an actor-critic architecture to label the importance of demonstrations. Coupled together with Hindsight Experience Replay (HER) [2], the trained agent learned how to stack up to six cubes, although it required the production and presentation of 100 demonstrations.

The question is, therefore, whether RL is capable of solving this stacking problem with the use of complex reward shaping and CL, as opposed to highly constrained and/or architecturally intricate scenarios or with the help of demonstrations. Examples of a complex problem solved using reward shaping and RL include Bellemare et al. [3] and Hu et al. [9]. The advantage that this architectural simplicity brings is a reduction of pipeline complexity for reproducibility and transferability, which can aid with safety constraints that can be added to the system later on.

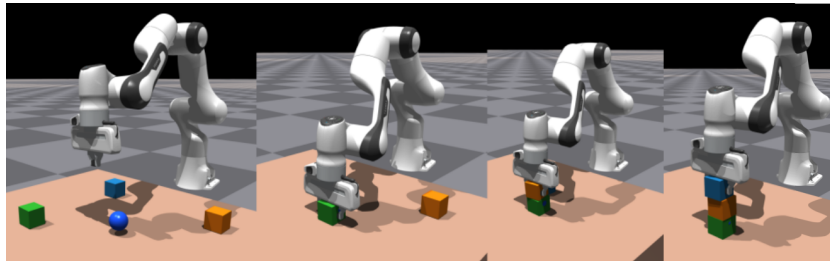
## 2.1 Reinforcement Learning

Throughout the paper, we will consider the goal-conditioned RL formalism [12] with a fully observable environment. The goals are individually sampled from a uniform distribution at the beginning of each episode. We utilise proximal policy optimization (PPO) [26] with the generalized advantage estimator (GAE) [25] to find the optimal policy. Our environment is built as an infinite-horizon setting with early stopping. We chose this setup because it significantly speeds up learning due to focusing the exploration around the starting position of the agent.

## 2.2 Curriculum Learning

Given a collection of tasks, each characterized by a reward function dependent on a target goal, a simple solution would involve attempting to learn all tasks simultaneously. This, unfortunately, is inefficient due to ignoring the dependencies between the tasks. Moreover, this solution assumes the agent is capable of learning any of these tasks in isolation, irrespective of the agent’s current capabilities. However, if the fraction of learnable tasks is small, this becomes computationally inefficient.

A better alternative is presented through the Curriculum Learning (CL) paradigm. A curriculum can be formed to narrow the distribution of tasks by training the agent only on tasks that it is capable of learning. Generally, it is unknown in advance which tasks can be learned given the current agent’s skills. However, there are environments, such as stacking of items, where prior knowledge can be utilised to select the right order for the tasks in which they are presented to the agent (i.e. gradually increase the number of items in the scene).



**Fig. 1.** We present our setup using a Franka Panda arm with 3 cubes in a simulation, created in Isaac Gym [18]. The image shows the intermediary goals achieved by a trained agent when stacking 3 objects.

In other words, CL is a training procedure that increases performance and/or training speed on a high-complexity task by decomposing it into a sequence of lower-difficulty sub-goals. Due to increased sample efficiency, the agent amasses knowledge faster and is able to generalize towards the complex task, reducing the exploration required if it were to train uninitialized.

The first application of CL in robotics dates back to 1994 in Sanger [23], where a parameter was used to change the difficulty of the environment in which the agent operates. Other influential work includes Schmidhuber [24], in which increasingly difficult sub-problems were identified and trained on. Florensa et al. [5] presented a novel and practical approach for sampling a start-state distribution that varies over the course of training, leading to an automatic curriculum of start-state distributions. Curriculum-based approaches with manually designed schedules have also been explored in supervised learning [7]. More recently, Forestier et al. [6] demonstrated how CL can be combined with intrinsic motivation to overcome certain robotics tasks.

For a more comprehensive overview of curriculum learning, please refer to Narvekar et al. [20] and Wang et al. [29].

### 2.3 Reward Shaping

Goal-oriented tasks are naturally described using sparse reward functions where the agent receives a positive reward only when achieving the task, and zero otherwise. Unfortunately, this makes the learning of such tasks difficult, due to the amount of guided exploration required to reach the goal even once. Most commonly [21], the process of reward shaping is used to overcome this issue. Reward shaping refers to modifying the original sparse reward function by incorporating domain knowledge in the form of additional reward.

Amodei et al. [1] and Grzes [8] observe that there are downsides to using a shaped reward. Without sufficient care, the changes brought to the reward function can mislead the agent towards a policy optimised for a local optimum introduced accidentally through the reward shaping. We explore these implications through an ablation study, which can be found in the Appendix.

### 3 Methods

The environment we use (Fig. 1) in our experiments consists of a Franka Panda (7 DOF) robotic arm, a table, and 3 boxes. The goal is to stack the objects on top of each other. The goal location and initial object positions are sampled randomly in every episode within the confines of the table. The table surface is split into a 5x5 grid, where each square represents a possible initial location for one of the cubes, thus preventing collisions from taking place during initialisation.

The initial joint positions are obtained by setting them to a hard-coded home position to which a small random perturbation (8%) is added:  $s_0 = s_{home} + 0.08 \cdot L \cdot \mathcal{U}[-1, 1]$ , where  $L$  represents the joint ranges. A total of 9 continuous actions are used, which include joint positions and gripper control. The observations are continuous and consist of joint position and velocity, end-effector position, object states, and goal position. More information can be found in the Appendix.

We implement the environment described above in Isaac Gym [18], an end-to-end high-performance robotics simulation platform. Throughout our experiments, we use the RL Games [17] (a highly-optimized GPU RL framework) implementation of Proximal Policy Optimization algorithm [26]. The Appendix contains more details regarding the hyperparameters used.

#### 3.1 Reward function

Similar to Bellemare et al. [3] and Hu et al. [9], the objective of the environment is encoded using a shaped reward function:  $r_t = \sum_{i \in \mathbb{S}} \lambda_i R_i$ , where  $\mathbb{S}$  represents the reward function set (see Table 1 for the  $\lambda$  values used). We start with the standard sparse reward for our environment, which rewards the agent when all the objects  $s_{obj}$  are placed at their associated goal positions  $s_g$  as shown in Eq. 1. We then add a dense reward as shown in Eq. 2, represented by the mean squared error between the current and target position of the objects. This error is summed across all objects.

$$r_{\text{sparse}} = \prod_{\text{obj} \in \mathbb{O}} \mathbb{1}\{s_{\text{obj}} \approx s_g \pm \epsilon\} \quad (1)$$

$$r_{\text{dense}} = - \sum_{\text{obj} \in \mathbb{O}} \|s_g - s_{\text{obj}}\|_2^2 \quad (2)$$

Additionally, we use a similar dense reward to guide the agent’s end effector  $s_{ee}$  toward the objects  $s_{obj}$  in the scene:  $r_{\text{guide}} = - \sum_{\text{obj} \in \mathbb{O}} \|s_{ee} - s_{\text{obj}}\|_2^2$ .

We introduce three additional signals for keeping the agent’s choice of action as safe and efficient as possible. The first reward signal punishes wasteful actions  $a$  and joint positions  $\theta$  that are far from the initial values:  $r_{\text{action}} = -\|a\|_2^2 - \|\theta - \theta^0\|_2^2$ . The second signal gives a high punishment whenever the end effector’s position on the Z axis  $s_Z^{ee}$  surpasses a threshold  $\Theta$ , equal to the table height:  $r_{\text{table}} = -\mathbb{1}\{s_Z^{ee} \leq \Theta\}$ . This prevents the agent from touching the table, which would otherwise cause problems when attempting to transfer the policy to a

real robot. Finally, the third signal rewards the agent for maintaining the end effector orientation  $\omega_{ee}$  close to the one it was initialised with  $\omega^0$ :  $r_{\text{orientation}} = -(\omega^0 * \hat{\omega}_{ee})$ . The orientations are represented in quaternions, and  $\hat{\omega}_{ee}$  symbolises the conjugate.

**Table 1.** Reward function set  $\mathbb{S}$

Signal	$\lambda$	Type
Goal achieved	150	$r_{\text{sparse}}$
One time bonus per sub-goal	150	$r_{\text{sparse}}$
Box to goal distance	5	$r_{\text{dense}}$
End-effector to box distance	5	$r_{\text{dense}}$
Action magnitude penalty	0.01	$r_{\text{action}}$
Touching table	5	$r_{\text{table}}$
Orientation error	0.1	$r_{\text{orientation}}$

### 3.2 Types of curriculum

In our experiments, we explore two curriculum setups. The switch between stages of a curriculum takes place automatically whenever the agent has reached an accuracy that surpasses a pre-set threshold. The accuracy is calculated as a moving average, updated each time one of the environments finishes an episode. In our experiments, we found that a threshold of 90% strikes the balance between time and optimality.

Our first curriculum setup consists of modifying the number of active objects in the scene. All three cubes are present on the table at all times, and their positions are used as part of the observation. The initial position of the cubes is selected from the 5x5 grid explained at the beginning of this section. This setup maintains an unchanged observation space throughout all stages of the curriculum.

The change, however, comes in the form of a gated reward function  $r_{\text{gated}} = \sum_i^{|\mathcal{O}|} r_{\text{sparse}} \cdot \mathbb{1}_{i \leq S}$ , such that at each stage  $S$ , the agent is encouraged to place one more object than in the previous stage. If a correctly placed object rewards the agent, we label them as *active* for the current curriculum stage. Otherwise, we call them *inactive*.

In our experiments, we gather trajectories for a number of steps equal to the horizon length  $H$ , allowing us to maintain a fixed batch size. The value of  $H$  can be tuned depending on the task, which introduces a new hyperparameter, requiring calibration. The only constraint we have is that  $H$  must be smaller than the episode length.

Our second curriculum setup builds on top of the gated reward function introduced above. It adjusts the episode and horizon length, such that, at each

stage, these values are updated based on the required time for solving the task with the active number of objects.

## 4 Input remapping

As a baseline, we start with what we call an absolute reward, meaning all rewards introduced in Section 3.1 are enabled. In this scenario, the episode only terminates when all present objects are positioned at their associated targets, or when the number of steps taken surpasses the episode length.

The results (Table 2) show that this baseline is incapable of learning anything meaningful during training. The gated reward curriculum, introduced in Section 3.2, represents our first proposed improvement. In this curriculum setup, as opposed to using the absolute reward, the episode ends when all *active* objects in the current curriculum stage are correctly placed.

Our first iteration establishes an order for the boxes to be placed in, similar to [16]. Moreover, we disable the observations associated with objects whose reward is gated, by replacing them with 0s. For example, during the first stage of the curriculum, only the red box contributes towards calculating the reward, and the observations of the other two inactive boxes are replaced with 0s before being passed to the agent.

In this iteration, the agent learns the sequence of actions that leads to the goal only with respect to the first box. When the second stage of the curriculum is reached, and the second object activates, the learned policy fails to complete the task associated with the first stage (i.e. place the red box at the bottom of the stack). This is because activating the second set of observations creates a strong shift in the input distribution to both policy and value function. Consequently, the generalization of knowledge from task 1 to task 2 is extremely dependent on the (often problematic) out-of-distribution generalization capability of the employed function approximators. This is similar to the problems encountered in the domain of continual learning [15].

Unfortunately, simply allowing the state of the inactive objects to contribute to the agent’s observation (instead of replacing them with 0s) does not have any beneficial effect on training. We explore a solution throughout the next sections.

### 4.1 Related Methods

The problem of input remapping has been studied on a different set of environments and confirmed by Tang and Ha [28]. Similarly, in Li et al. [16], the authors achieved high accuracy in stacking up to 6 cubes, by using a curriculum and an architecture of GNNs with attention. GNNs provide input permutation invariance, thus overcoming the problem of input remapping. Our work differs through architectural simplicity and the insights discovered that enabled the learning to take place.

Alternative solutions include the use of elastic weights consolidation [13], or of an auto-encoder [22]. The first method identifies the weights that are essential

to previously learned tasks and reduces further changes, such that the behaviour is preserved. This is achieved by freezing the node weights that have contributed most to learning that task. The second method overcomes the remapping issue by translating the states into a latent representation, formed using a variational auto-encoder, that is invariant to the permutation of inputs. Both of these solutions require architectural changes and may be explored in more detail in a separate study.

Our hypothesis is that the lack of permutation invariance of standard neural networks is one of the reasons why vanilla RL methods were not capable of solving the problem of stacking multiple objects, particularly in a curriculum setting. Without addressing this issue first, the curriculum of sub-tasks brings almost no benefits to the training, since the trajectory of rewarding actions has to be learned from scratch whenever a new object is introduced.

## 4.2 Results

We identified a solution to the input remapping issue presented above that requires no architectural changes, however, all objects must be present from the start. This is overcome by crafting the reward function in a way that guides the agent to interact with all the boxes even in the first stage of the curriculum. To achieve this, every time the episode is reset, a permutation of boxes is chosen randomly, such that a different stacking order is required for achieving the goal. As a result, all input weights associated with objects are trained equally.

Therefore, our final iteration for this first curriculum setup consists of enabling the observations for all objects, where the order for the gated reward used is permuted for each episode. An episode ends when either enough timesteps have passed, or all the *active* objects in the current stage have been positioned correctly.

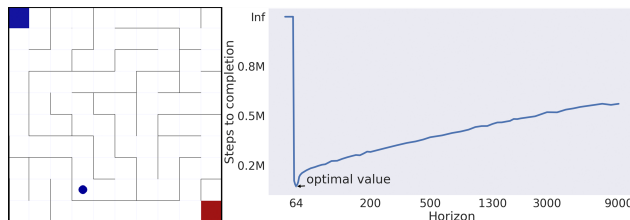
Lastly, to better understand the performance boost brought by the curriculum, we used a staggered reward as a baseline, similar to the one described in [3]. Here, we still make use of the gated reward function, however, the episode only ends when all three objects are correctly positioned. See Table 2 for a comparison of our methods.

The curriculum run that achieved the best results was trained for approximately 21.5 billion steps, which is equivalent to 20 hours on our workstation (see Appendix for hardware details).

**Table 2.** Results of the various iterations of the first curriculum

Name	Reward	Input	Ep. stop	Obj. order	Acc.	Ep. L.	Steps	Time
Absolute	normal	all	all obj	fixed	0%	-	30B	45h
Curriculum 1a	gated	padded	active obj	fixed	0%	-	30B	28h
Curriculum 1b	gated	all	active obj	fixed	0%	-	30B	28h
Staggered	gated	all	all obj	fixed	81%	90	30B	31h
Curriculum 1c	gated	all	active obj	permutation	<b>90%</b>	<b>82</b>	<b>21.5B</b>	<b>20h</b>





**Fig. 2.** Left: 10x10 maze environment with discrete action space using sparse rewards provided when the agent reaches the goal (red square). Right: Training steps required to learn a policy capable of solving the maze environment over a range of values for horizon length. Providing more time in each episode seems to have diminishing returns for the overall training time.

## 5 Horizon length

PPO [26] is an on-policy online RL method, meaning that updates take place after each episode terminates using the gathered transitions. The algorithm was designed with parallelism in mind, where each agent collects the same amount of  $H$  timesteps before each update. Further studies [27] into the efficiency of parallelization considered only score-based environments, such as the Pong and Breakout games, but failed to analyse the importance of the horizon window when dealing with goal-oriented scenarios. In their bipedal locomotion environment, the agent can make incremental improvements as long as it is not in an unrecoverable state (e.g. fallen). However, we show in this section that in a goal-oriented task, the size of the learning horizon  $H$  is an important hyperparameter that offers a trade-off between optimality and robustness. While this may seem obvious to the reader, there have been no published results in this area.

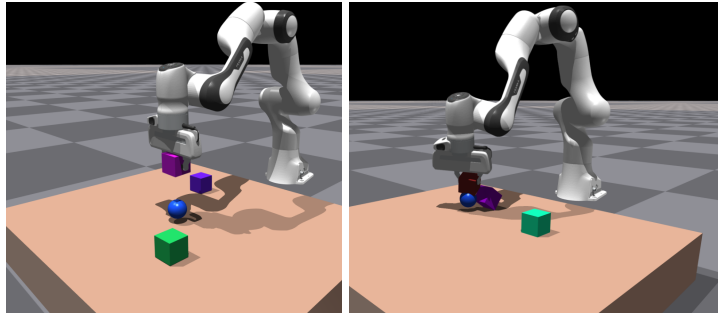
Our second curriculum setup consists of increasing the horizon and episode length at each stage, such that the agent has more time to explore and solve the task before being reset. Our hypothesis is that choosing a value of  $H$  close to optimal for solving the task at each stage will decrease the training time.

Ideally, we want to select the optimal  $H$  value for all tasks. Unfortunately, this value is hard to determine without solving the task beforehand, especially for long-horizon tasks requiring multiple checkpoints to be reached, such as object stacking. Therefore, we will next study the effect of the horizon length on convergence speed using a toy example (see Fig. 2).

The results show that there is an optimal value and any further increase provides diminishing returns, leading to a less-than-optimal training time. Intuitively, if  $H$  drops below the required steps for solving the task, the trajectories gathered do not encapsulate the solution, hence making the problem impossible to learn. Knowing this, we empirically estimated a base horizon length  $H_{\text{base}}$  of 50 steps per *active* object. This is a good trade-off between the risk associated with having a too-small horizon length, and the sub-optimality of a too-large value.

The transition between stages in this setup is identical to the first curriculum, presented in Section 3.2. The horizon length is updated as follows:  $H = H_{\text{base}} * S$ , where  $S$  represents the stage number.

### 5.1 Reward hacking



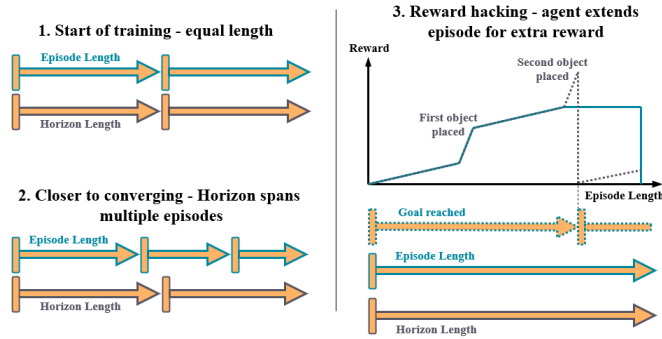
**Fig. 3.** The first image shows an example of reward hacking where the agent is using a single cube to collect all three one-time bonus rewards. The second image shows the robotic arm pushing aside the first correctly placed box, therefore postponing the end of the episode.

Throughout our experiments concerning the second curriculum, we came across a couple of reward-hacking scenarios (see Fig. 3) directly linked to the tuning of the horizon length. This section is dedicated to explaining these instances.

In Stooke and Abbeel [27], each parallel agent has access to its own simulator where experiences are gathered asynchronously from the learner and resets can happen in sync with data acquisition. Isaac Gym [18] reaches a high number of parallel environments by creating them all using the same simulator. As a consequence of maintaining a fixed batch size, the training takes place after the number of steps dictated by the horizon length has been gathered across all environments.

If the episode and horizon lengths are equal and the agent cannot achieve the goal (i.e. at the beginning of training), the asynchronous experience gathering is equivalent to the synchronous method. However, when the agent starts achieving the goal with higher accuracy, the environment resets are out of sync, therefore the horizon window can span multiple episodes (see Fig. 4).

Within the second setup discussed in this section, the agent finds a reward hack during the second phase of the curriculum. A policy is found that can generate the most reward in a window by placing the first box, but not ending the episode with the second box. Instead, it displaces the first box using the second, refusing to finish the episode early such that it circumvents the early negative reward of a fresh episode.

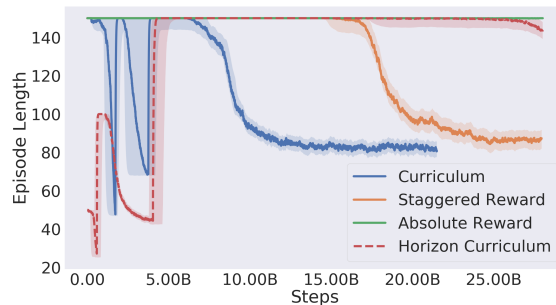


**Fig. 4.** 1. Initially, the episode and horizon lengths are equal. 2. Once the agent has learned part of the task, the episode lengths are smaller than the horizon. 3. Reward hacking scenario with 2 active objects where the agent maximizes the return over the horizon length, rather than episode’s duration. Dotted line shows desired behaviour, however, the agent extends the episode instead.

In order to incentivise the agent to finish the episode, we increase the bonus reward for achieving any sub-goal, including the overall environment goal. This, however, requires a balance, since, if this number is set too high, the agent converges on a local optimum, thereby throwing the objects toward the target locations, without actually building a tower.

The second curriculum implementation takes considerably longer to train when compared to the baseline of the staggered reward, introduced in Section 4.2. Our hypothesis is that an increase in horizon length leads to a drastic change in the expected reward for each state. The returns are approximated using the GAE [25] method, and the value of a state is estimated based on the amount of reward that can be gathered until the end of the episode, starting from that state. Since the agent can now explore more states, due to the increase in maximum timesteps, the expected reward also increases in magnitude. The agent cannot progress on learning the requirements for the new sub-goal until the value function is updated based on the new horizon length. This additional time required is what we believe is slowing down training.

To test this, we consider an alternative setup, where  $H$  is set to a close-to-optimum value for the last stage of the curriculum directly from the beginning. This means that the early phases of the curriculum will take longer to train due to a non-optimal  $H$  value for that stage. However, as shown in Fig. 5, this change speeds up training in the long run. Therefore, our initial hypothesis that a curriculum of horizon length speeds up training has been disproved. Table 3 encapsulates the ablation study that compares different setups for episode and horizon length over 3 billion steps and with 2 *active* objects.



**Fig. 5.** Comparison of curriculum setups against the chosen baselines. Curriculum - Curriculum 1c with  $H = 150$ . Staggered and Absolute reward - as described in Table 2 with  $H = 150$ . Horizon Curriculum - second curriculum setup with  $H = [50, 100, 150]$ . Curriculum (blue) is the only run achieving 90% accuracy.

**Table 3.** Second curriculum ablation results with 2 active objects

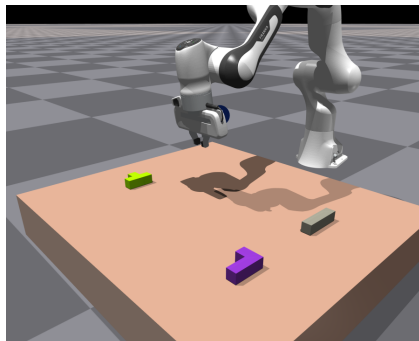
Max Ep. Length	Horizon Length	Accuracy	Average Ep. Length
100	100	90%	58
50-100	50-100	67%	62
100	10	0%	100
100	50-100	71%	64

## 6 Increasing complexity

Additionally, we test the first curriculum setup on a more challenging environment (Fig. 6). In this setting, we replace the cubes with three pieces from the set of five tetrominoes (L, T, and I). The increase in complexity comes from the need to learn new different grasping methods associated with each object. To account for this, we increase the network capacity by one layer. In our experiments, the items just need to be stacked in a specified position, although more complexity can be introduced in later experiments, where a certain orientation is required.

We choose the objects to have the same volume as the previously used cubes, however, this led to smaller grasping areas. Initially, the agent fails to learn a good policy within the same time frame. To improve the results, we add a new reward signal  $r_{\text{grasp}} = \mathbb{1}\{|s_{\text{ee}} - s_{\text{obj}}| \leq \epsilon\}$  to assist in guiding the grasping. We find a  $\lambda_{\text{grasp}} = 2$  to be most effective.

In a separate test, we increase the complexity further by providing a random initial orientation to each object. To allow the agent to adjust to this change, we concatenate the end-effector’s rotation (quaternions) to the observations. The results in Table. 4 show that curriculum learning, combined with reward shaping, enables learning even in this complex scenario.



**Fig. 6.** Environment consisting of three pieces from the set of five tetrominoes (L, T, and I), each being initialised with a fixed orientation. The different shapes increase complexity for the agent.

**Table 4.** Results in the higher complexity environment

Name	Task	$\lambda_{\text{grasp}}$	Rand. Or.	Acc.	Steps
Cubes	cubes	0	no	90%	16B
Complex tetrominoes	2	no		80%	16B
RandOr	tetrominoes	2	yes	59%	16B

## 7 Limitations

The trained agents are currently unable to cope with the situation where the tower of stacked objects is knocked down. Instead of trying to reposition the previous objects, the agent continues as though nothing happened. We leave this investigation for future work.

Another limitation of the system presented is that the agent is trained on a maximum of  $N$  objects. In order to have  $N + 1$  objects, the training must restart from scratch. In further work, we intend to deal with this scalability issue.

## 8 Conclusion

In this work, we investigated the problem of applying curriculum learning to a goal-oriented environment, such as the stacking of boxes. To accomplish this, we trained an RL agent using a shaped reward function and coupled it with multiple variations of curricula architectures. We found that a simple architecture combined with a curriculum setting can achieve the desired behaviour, whereas vanilla RL fails to complete the task. We provided solutions for a number of reward-hacking scenarios identified and showcased results obtained in a higher-complexity environment. Altogether, these findings establish a strong foundation for future research on CL applied to goal-oriented tasks.

## Bibliography

- [1] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mané. Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*, 2016.
- [2] M. Andrychowicz, F. Wolski, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, O. Pieter Abbeel, and W. Zaremba. Hindsight experience replay. *Advances in neural information processing systems*, 30, 2017.
- [3] M. G. Bellemare, S. Candido, P. S. Castro, J. Gong, M. C. Machado, S. Moitra, S. S. Ponda, and Z. Wang. Autonomous navigation of stratospheric balloons using reinforcement learning. *Nature*, 588(7836):77–82, 2020.
- [4] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- [5] C. Florensa, D. Held, M. Wulfmeier, M. Zhang, and P. Abbeel. Reverse curriculum generation for reinforcement learning. In S. Levine, V. Vanhoucke, and K. Goldberg, editors, *Proceedings of the 1st Annual Conference on Robot Learning*, volume 78 of *Proceedings of Machine Learning Research*, pages 482–495. PMLR, 13–15 Nov 2017. URL <https://proceedings.mlr.press/v78/florensa17a.html>.
- [6] S. Forestier, R. Portelas, Y. Mollard, and P.-Y. Oudeyer. Intrinsically Motivated Goal Exploration Processes with Automatic Curriculum Learning. To be published at *Journal of Machine Learning Research* 2022, Nov. 2021.
- [7] A. Graves, M. G. Bellemare, J. Menick, R. Munos, and K. Kavukcuoglu. Automated curriculum learning for neural networks. In *international conference on machine learning*, pages 1311–1320. PMLR, 2017.
- [8] M. Grzes. Reward shaping in episodic reinforcement learning. 2017.
- [9] Y. Hu, W. Wang, H. Jia, Y. Wang, Y. Chen, J. Hao, F. Wu, and C. Fan. Learning to utilize shaping rewards: A new approach of reward shaping. *Advances in Neural Information Processing Systems*, 33:15931–15941, 2020.
- [10] A. Hundt, B. Killeen, N. Greene, H. Wu, H. Kwon, C. Paxton, and G. D. Hager. 'good robot!': Efficient reinforcement learning for multi-step visual tasks with sim to real transfer. *IEEE Robotics and Automation Letters*, 5(4):6724–6731, 2020.
- [11] R. Jeong, Y. Aytar, D. Khosid, Y. Zhou, J. Kay, T. Lampe, K. Bousmalis, and F. Nori. Self-supervised sim-to-real adaptation for visual robotic manipulation. In *2020 IEEE international Conference on Robotics and Automation*, pages 2718–2724. IEEE, 2020.
- [12] L. P. Kaelbling. Learning to achieve goals. In *IJCAI*, volume 2, pages 1094–8. Citeseer, 1993.
- [13] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, et al. Over-

- coming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.
- [14] A. X. Lee, C. M. Devin, Y. Zhou, T. Lampe, K. Bousmalis, J. T. Springenberg, A. Byravan, A. Abdolmaleki, N. Gileadi, D. Khosid, C. Fantacci, J. E. Chen, A. Raju, R. Jeong, M. Neunert, A. Laurens, S. Saliceti, F. Casarini, M. Riedmiller, r. hadsell, and F. Nori. Beyond pick-and-place: Tackling robotic stacking of diverse shapes. In *Proceedings of the 5th Conference on Robot Learning*, volume 164 of *Proceedings of Machine Learning Research*, pages 1089–1131, 08–11 Nov 2022.
- [15] T. Lesort, V. Lomonaco, A. Stoian, D. Maltoni, D. Filliat, and N. Diaz-Rodriguez. Continual learning for robotics: Definition, framework, learning strategies, opportunities and challenges. *Information fusion*, 58:52–68, 2020.
- [16] R. Li, A. Jabri, T. Darrell, and P. Agrawal. Towards practical multi-object manipulation using relational reinforcement learning. In *2020 IEEE International Conference on Robotics and Automation*, pages 4051–4058. IEEE, 2020.
- [17] D. Makoviichuk and V. Makoviychuk. RL-games: A highperformance framework for reinforcement learning, 2022. URL <https://github.com/Denys88/rl-games>.
- [18] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa, et al. Isaac gym: High performance gpu-based physics simulation for robot learning. *arXiv preprint arXiv:2108.10470*, 2021.
- [19] A. Nair, B. McGrew, M. Andrychowicz, W. Zaremba, and P. Abbeel. Overcoming exploration in reinforcement learning with demonstrations. In *2018 IEEE international Conference on Robotics and Automation*, pages 6292–6299. IEEE, 2018.
- [20] S. Narvekar, B. Peng, M. Leonetti, J. Sinapov, M. E. Taylor, and P. Stone. Curriculum learning for reinforcement learning domains: A framework and survey. *arXiv preprint arXiv:2003.04960*, 2020.
- [21] A. Y. Ng, D. Harada, and S. Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In *Icml*, volume 99, pages 278–287, 1999.
- [22] A. Raffin, A. Hill, R. Traore, T. Lesort, N. Diaz-Rodriguez, and D. Filliat. Decoupling feature extraction from policy learning: assessing benefits of state representation learning in goal based robotics. *arXiv preprint arXiv:1901.08651*, 2019.
- [23] T. D. Sanger. Neural network learning control of robot manipulators using gradually increasing task difficulty. *IEEE transactions on Robotics and Automation*, 10(3):323–333, 1994.
- [24] J. Schmidhuber. Powerplay: Training an increasingly general problem solver by continually searching for the simplest still unsolvable problem. *Frontiers in psychology*, 4:313, 2013.
- [25] J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*, 2015.

- [26] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [27] A. Stooke and P. Abbeel. Accelerated methods for deep reinforcement learning. *arXiv preprint arXiv:1803.02811*, 2018.
- [28] Y. Tang and D. Ha. The sensory neuron as a transformer: Permutation-invariant neural networks for reinforcement learning. *Advances in Neural Information Processing Systems*, 34, 2021.
- [29] X. Wang, Y. Chen, and W. Zhu. A survey on curriculum learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.



## A APPENDIX

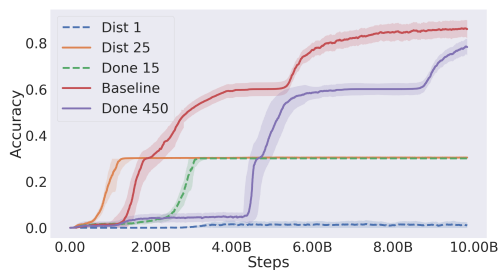
Code available post review. Hardware used: 2x Nvidia GeForce RTX2080ti.

**Table A1.** Observation space

Observation	Size Info	
Joint Positions	9	Scaled between -1 and 1
Joint velocities	9	
End-effector position	3	
Object states	39	Position, velocity and orientation
Goal position	3	
Object id	1	Used for staggered reward
Total	64	

**Table A2.** Hyperparameters

Hyperparameter	Value
Algorithm	PPO
Separate Actor/Critic	False
MLP size	[256, 256, 256]
Fixed $\sigma$	False
Activations	ELU
Learning rate ( $\eta$ )	5e-4
Learning rate type	Fixed
Discount factor $\gamma$	0.99
GAE $\tau$	0.95
KL threshold	0.008
Nr environments	8192 - 16384
Minibatch size	32768
Horizon/Ep length	150



**Fig. 7.** Comparison of various  $\lambda$  magnitudes affecting distance-based dense and sparse rewards. 'Done 15' and 'Dist 25' successfully solve stage 1 but fail to progress further. 'Dist 1' fails to find any optimal policy.

**Reward function ablation**

We have conducted an ablation study for the values of  $\lambda$  used in adjusting our reward function. The values adjusted include the 'Box to Goal' and 'End-effector to Box' distances and the 'Goal achieved' and 'One time bonus per sub-goal' sparse rewards. These are represented in Figure 7 as 'Dist' and 'Done', respectively. The baseline consists of  $\lambda_{\text{dist}} = 5$  and  $\lambda_{\text{done}} = 150$ , which represent the values we have used throughout our experiments.