



HAL
open science

GRIMGEP: Learning Progress for Robust Goal Sampling in Visual Deep Reinforcement Learning

Grgur Kovač, Adrien Laversanne-Finot, Pierre-Yves Oudeyer

► **To cite this version:**

Grgur Kovač, Adrien Laversanne-Finot, Pierre-Yves Oudeyer. GRIMGEP: Learning Progress for Robust Goal Sampling in Visual Deep Reinforcement Learning. IEEE Transactions on Cognitive and Developmental Systems, 2022, 10.1109/TCDS.2022.3216911 . hal-03121352v2

HAL Id: hal-03121352

<https://inria.hal.science/hal-03121352v2>

Submitted on 22 Nov 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

GRIMGEP: Learning Progress for Robust Goal Sampling in Visual Deep Reinforcement Learning

Grgur Kovač

Flowers Team

INRIA (FR)

grgur.kovac@inria.fr

Adrien Laversanne-Finot

Flowers Team

INRIA(FR)

adrien.laversanne-finot@inria.fr

Pierre-Yves Oudeyer

Flowers Team

INRIA(FR)

pierre-yves.oudeyer@inria.fr

Abstract

Designing agents, capable of learning autonomously a wide range of skills is critical in order to increase the scope of reinforcement learning. It will both increase the diversity of learned skills and reduce the burden of manually designing reward functions for each skill. Self-supervised agents, setting their own goals, and trying to maximize the diversity of those goals have shown great promise towards this end. However, a currently known limitation of agents trying to maximize the diversity of sampled goals is that they tend to get attracted to noise or more generally to parts of the environments that cannot be controlled (distractors). When agents have access to predefined goal features or expert knowledge, absolute Learning Progress (ALP) provides a way to distinguish between regions that can be controlled and those that cannot. However, those methods often fall short when the agents are only provided with raw sensory inputs such as images. In this work we extend those concepts to unsupervised image-based goal exploration. We propose a framework that allows agents to autonomously identify and ignore noisy distracting regions while searching for novelty in the learnable regions to both improve overall performance and avoid catastrophic forgetting. Our framework can be combined with any state-of-the-art novelty seeking goal exploration approaches. We construct a rich 3D image based environment with distractors. Experiments on this environment show that agents using our framework successfully identify interesting regions of the environment, resulting in drastically improved performances. The source code is available at <https://sites.google.com/view/grimgep>.

1 Introduction

Recent work in reinforcement learning has shown that robots can learn complex individual skills such as grasping [Kalashnikov et al., 2018], locomotion [Haarnoja et al., 2018a, Yang et al., 2020], and manipulation tasks [Gu et al., 2017]. However, the standard paradigm in reinforcement learning often involves learning policies specific to each task. This imposes an engineering burden as a reward function has to be manually designed for each task. On the contrary, humans are capable to autonomously learn a wide variety of skills by setting their own goal and practicing them, with little to no supervision. They naturally move from simple tasks in the early infancy to more complex tasks as they grow. Designing autonomous learning agents, capable of setting their own goals, similarly to how human learn is thus a promising avenue of research. Such agents would autonomously discover

general purpose policies, that could later be used for specific purpose, reducing the need for outside intervention and manual engineering.

In the absence of a specific objective, an efficient strategy in order to learn a general purpose policy is to learn to reach as many states as possible [Nair et al., 2018, Pong et al., 2020, Nasiriany et al., 2021, Laversanne-Finot et al., 2018, Colas et al., 2019]. Combining a goal reaching policy with an exploration mechanism, even in its simplest form, can allow solving otherwise impossible environments. For example, GO-Explore [Ecoffet et al., 2021] showed that when goal-reaching is assumed solved by expert knowledge (manually setting the environment state to the goal state), performing random actions from interesting goals states allows to solve hard exploration problems.

This form of goal exploration involves two challenges: how to learn a goal-reaching policy and how to explore as many novel goals as possible. This first challenge is solved by using recent reinforcement learning algorithms [Andrychowicz et al., 2017, Haarnoja et al., 2018b] together with good goal representations [Nair et al., 2018, Warde-Farley et al., 2018, Laversanne-Finot et al., 2018]. In order to solve the second issue, many approaches try to maximize the diversity of sampled goals, for example by using importance sampling techniques [Pong et al., 2020] in order to sample states uniformly from the underlying distribution.

While optimizing the diversity of sampled goals has been shown to enable robots to solve complex manipulation tasks [Pong et al., 2020], it falls short in environments containing unpredictable distractors. This problem is often referred to as the noisy TV problem as an agent aiming at maximizing novelty would get attracted by a TV displaying a sequence of continually novel and unpredictable images [Kaplan and Oudeyer, 2007, Burda et al., 2019].

One way to solve this problem is to explore parts of the environment with high learning progress [Schmidhuber, 1991, Oudeyer et al., 2007], which is called competence progress when it measures progress towards self-generated goals [Baranes and Oudeyer, 2013, Forestier et al., 2017]. For example, SAGG-RIAC [Baranes and Oudeyer, 2013] gradually clusters the environment into sub-regions of different learning progress and samples goals in the more promising regions (e.g. the regions with high learning progress). This approach has been shown to enable learning of inverse models in high-dimensional robots. In other approaches, the goal space is organized along several modules encoding for different types of goals: in that case goal sampling is hierarchical [Forestier et al., 2017, Colas et al., 2019]. Another interesting aspect of learning progress is that it allows avoiding catastrophic forgetting [Colas et al., 2019] by identifying skills that were previously mastered, but are not mastered anymore. Using the learning progress to sample goals in image based environments is challenging due to the high dimensionality of the observations and as only been done in simple settings [Laversanne-Finot et al., 2018].

In this paper we provide a solution to action-induced distractors in image-based environments for RL-based goal exploration algorithms. The general idea is to cluster the environment into different parts for which the agent can then estimate the associated learning progress. From a high level perspective our method can be viewed as a way to learn a prior over possible goals. The goal of this prior is to detect which parts of the environment should be explored. This prior can then be combined with any goal exploration algorithm in order to guide exploration. We experimentally demonstrate that our framework improves the overall performances of various novelty-based goal exploration algorithms and autonomously focuses the exploration on interesting learnable regions of the environment.

Contributions The main contribution of this work is to study how learning progress can be estimated in image-based environments and how it can be used to guide the exploration of existing novelty-based goal exploration algorithms in Deep RL. The main contributions are:

- We introduce the GRIMGEP framework that allows estimating learning progress in image-based environments for goal sampling in goal-conditioned Deep RL
- We combine it with existing goal exploration algorithms, and show that it improves the performances and avoids catastrophic forgetting of these algorithms in the presence of distractors.
- We provide an easily customizable 3D image-based environment constructed for studying noisy distractors in goal driven exploration tasks.

2 Related work

Goal exploration algorithms Intrinsically Motivated Goal Exploration Processes (IMGEPs) are a class of exploration algorithms that explore by repeatedly setting goals for themselves that they then try to achieve. They exist in two forms: population-based IMGEPs leveraging non-parametric models (e.g. Baranes and Oudeyer [2013], Forestier et al. [2017]), and, as we study in this paper, goal-conditioned RL IMGEPs leveraging Deep RL techniques (e.g. Nair et al. [2018], Pong et al. [2020], Colas et al. [2019], Nasiriany et al. [2021]). Combined with an efficient curiosity mechanism for sampling goals, this approach has been shown to enable high-dimensional robots to learn very efficiently locomotion skills [Baranes and Oudeyer, 2013], manipulation of objects [Rolf et al., 2010, Nguyen and Oudeyer, 2014, Nair et al., 2018, Pong et al., 2020] or tool use [Forestier et al., 2017]. Methods using absolute learning progress [Baranes and Oudeyer, 2013] to drive goal sampling were shown to scale up to real world environments with potentially many forms of distractors, including action-induced distractors [Forestier et al., 2017, Colas et al., 2019]. However, these works relied on abstract hand-defined goal and state spaces (e.g. based on object positions and velocities). An exception is Laversanne-Finot et al. [2018], which used learning progress to sample goals in learned latent spaces, but this relied on population-based learners and required offline pre-training. Here, we adapt the learning progress approach to Deep RL goal-exploring agents that perceive their environments through pixels, and study their properties in this context.

Recently some approaches studied goal exploration in the image-based goal-conditioned Deep RL framework. For instance, Nair et al. [2018] learns a goal-conditioned policy on top of a learned embedding of the environment. Nasiriany et al. [2021] proposes a richer task representation in the form of a goal distribution. Warde-Farley et al. [2018] learns a goal conditioned policy by maximizing the mutual information between the goal state and the achieved state. While the goal policy achieves a good performance, the heuristic used for goal sampling is very simple and was not shown to scale to large environments with distractors. Pong et al. [2020] improves upon Nair et al. [2018] by designing a mechanism that incentivizes goal sampling to focus on exploring the frontier of the distribution of known goals, implementing a form of novelty search Lehman and Stanley [2011] (see also Eysenbach et al. [2018] for a related approach). However, such methods can be naturally attracted by distractors that generate unpredictable or uncontrollable novel observations (either due to randomness, lack of cognitive capacity or lack of access to hidden information). Here, we experimentally identify this limit, and show that the learning progress based goal sampling mechanism we introduce can be used as a wrapper around these algorithms to enable them to avoid irrelevant or unlearnable regions of the environment, while keeping their efficient novelty-based exploration in learnable regions.

Curiosity-driven Deep RL A common trend to improve exploration in the classical reinforcement learning setting with sparse external rewards has been to supplement the task reward with intrinsic reward [Bellemare et al., 2016, Pathak et al., 2017, Burda et al., 2018a, Raileanu and Rocktäschel, 2020]. Approaches considering image-based low-level perception have used intrinsic rewards measuring various forms of novelty, based on counts [Bellemare et al., 2016] or prediction errors [Pathak et al., 2017, Burda et al., 2018b]. While these approaches can give impressive results, they have fallen short on simple environments in the presence of a distractor that is partially controlled by one of the agent’s actions. Kim et al. [2020] studied the use of learning progress as an intrinsic reward to enable Deep RL agents to be robust to distractors. However, this approach relied on high-level disentangled state representations and did not include the notion of goal exploration.

3 Problem definition

Open-ended unsupervised exploration in image based environments

In the problem of open-ended unsupervised exploration the agent is tasked with acquiring a diversity of skills while having no access to the internal dynamics of the environment or expert knowledge. The only information the agent receives from the environment is in the form of observations that it collects. The agent must discover new states and learn how to reach them at the same time, often in the form of a goal policy.

This problem is particularly challenging when the agent has only access to observations in the form of images for two main reasons. First, since no measure of similarity between a state image and the goal image is given, the agent has to autonomously train the goal-conditioned reward function

in an unsupervised manner. Second, since no information is given about which images represent feasible states in the environment, a mechanism must be constructed to autonomously create and sample plausible goals.

Algorithm 1: Intrinsically Motivated Goal Exploration Process in the Deep RL setting (blue lines correspond to GRIMGEP additions)

```

π.init()  $\mathcal{R}$ .init() // goal-conditioned policy and the reward function
encountered_states = [] performance_history = []
// random acting rollouts to fill the replay buffer
for  $N_{warmup}$  do
    traj = env.random_rollout()
    encountered_states.add_states(traj)
// exploration
for ep in  $N_{epochs}$  do
    if ep > start_exploration then
        goal_intrinsic_rewards = compute_intrinsic_rewards(encountered_states) // novelty-based
        clusters = clustering_fn(performance_history)
        per_cluster_alp = estimate_ALPs(clusters, performance_history)
        prior = construct_prior(clusters, per_cluster_alp) // sample a cluster
        goal_intrinsic_rewards *= prior // by eq. 3
    else
        goal_intrinsic_rewards = [1, 1, ..., 1] // at the beginning sample uniformly
        goal = prioritized_sample_goal(encountered_states, goal_intrinsic_rewards)
        trajectory = env.policy_rollout(π, goal)
        encountered_states.add_states(trajectory)
        performance_history.add_pair((goal, trajectory.last_state))

     $\mathcal{R}$ .train() π.train() // fit VAE and policy
    clustering_fn.train(encountered_states) // fit PCA and GMM

```

Intrinsically motivated goal exploration

Intrinsically motivated goal exploration processes (IMGEP) are a natural way to address this problem. IMGEPs are a family of approaches that explore by repeatedly setting goals and then trying to achieve them. Those approaches usually consist of two parts: a goal reaching policy, and a goal sampling mechanism. Having a good policy ensures that goals are reached efficiently, and the goal sampling mechanism serves to discover new goals. The combination of these two elements ensures that the agent explores the environment efficiently. In this work we focus on the goal sampling mechanism.

In the case of image-based environments, goals cannot be sampled directly from the image space as such images are unlikely to represent images feasible in the environment. A simple solution, is to sample goals from the history of observations. However, in many cases, this history will not be filled uniformly from observations of the environment. For example, it will often be biased towards observations that are close to the starting position of the agent. Thus, sampling uniformly goals from that history will not provide very diverse goals. Therefore, in order to find interesting goals, current state-of-the-art methods rely on sampling methods based on the novelty of the visited states to maximize the diversity of sampled goals. A pseudocode of such an algorithm is shown by black lines in Algorithm 1.

However, those novelty-seeking approaches fall short in environments with distractors, i.e. where unpredictable (possibly structured) images can be perceived in certain situations (like a TV showing a flow of continuously novel images). Such distractors are problematic as a goal depicting a specific randomly occurring state is infeasible. Furthermore, as they produce a large amount of novelty, current approaches are attracted to them and oversample infeasible goals resulting in reduced performance and catastrophic forgetting.

IMGEPs using Absolute Learning Progress (ALP) instead of novelty have been shown to address both of those issues in simpler (not image-based) environments. In this paper, we present the **GRIMGEP** approach that extends the ALP based exploration to unsupervised image-based exploration. GRIMGEP uses ALP estimates to detect learnable regions of the goal space and then uses

current novelty-based approaches only inside these regions. The extensions introduced by GRIMGEP are shown by blue lines in Algorithm 1.

Evaluation

In order to evaluate such goal exploration algorithm, the evaluation objective used by the experimenter consists in measuring the performance of the learned goal-reaching policy over a test set of goals which represents well the diversity of learnable skills in the chosen environment (note the distribution of goals in the test set is unknown to the learner):

$$\operatorname{argmax}_{\theta} \int_{g \sim \mathcal{G}} f(g, \tau_{\pi_{\theta}, g}) dg \quad (1)$$

Where \mathcal{G} is the goal space, $\tau_{\pi_{\theta}, g}$ the trajectory resulting from following the goal-conditioned policy π , parameterized by θ , and aiming for goal g . f is the goal fulfillment evaluation function specifying to what extent was the goal g reached in episode $\tau_{\pi_{\theta}, g}$. This function is also defined by the experimenter for the purpose of evaluation, and the agent has no access to it during training.

In our case of image-based goal exploration, f is a measure of similarity between the goal and the final state i.e. to what extent are the same objects visible on the two images. Since it is unfeasible to integrate over the whole image space, a fixed set of testing goals (test set) is used for evaluation. For further details about the evaluation and the evaluation metric refer to sections 6.1, and appendix C.

4 Current novelty-based exploration approaches

In this work, we consider two representative novelty seeking goal exploration approaches: Skewfit [Pong et al., 2020] and *CountBased*. Both Skewfit and Countbased train a generative model (β -VAE) online on the encountered states. They use the negative $L2$ distance between the goal and the state in this latent space as the reward for training the goal-conditioned policy. Both sample goals from a set of all the encountered states, however the way in which they prioritize goal sampling is different. This same prioritization method is also used for sampling the data to train the generative model.

Skewfit [Pong et al., 2020] was shown to outperform a number of baselines for unsupervised goal exploration in environments without noisy distractors. It samples goals prioritized by how novel they appear to the current generative model (VAE) i.e. the probability of sampling a goal is inversely proportional to the probability of that goal under the generative model’s distribution. For a more detailed description of Skewfit refer to Appendix A.1.

CountBased’s goal sampling is inspired by how Go-Explore samples cells [Ecoffet et al., 2021]. The states are downsampled to a 3x3 image, and quantized to 4 values (4^3 different colors). All the encountered states are counted as their quantized versions and the goal sampling biased accordingly, i.e. the probability of sampling a goal is inversely proportional to the number times its quantized version was encountered. For a more detailed description of CountBased refer to appendix A.2.

5 GRIMGEP

We present the **Goal Regions guided Intrinsically Motivated Goal Exploration Process (GRIMGEP)** framework. The GRIMGEP framework can be combined with any goal exploration algorithm, improving it by allowing the agent to identify and to sample goals from the interesting and learnable regions of the goal space.

GRIMGEP extends the ALP approach to unsupervised image-based exploration by fitting a clustering function on the storage of encountered states, applying this clustering function to a separate history of (attempted goal, achieved last state) pairs, estimating ALP for each of the clusters using this history, and sampling a cluster based on these ALP estimates. The wrapped IMGEP (which can be novelty based) is then used to sample a goal from the storage off all the encountered states, but constrained to those goals that the clustering function would assign to the sampled cluster. This pipeline is repeated for every sampled goal.

The motivation behind this scheme is that clusters corresponding to distracting unlearnable regions will have small ALP (performance doesn’t change) and will be sampled rarely, while regions that are currently being learned or forgotten (performance rising or falling) will be sampled more.




Figure 1: Goal sampling procedure in the GRIMGEP framework. 1) First, the **Clustering component** clusters the goal space i.e. fits a clustering function. 2) The absolute learning progress (ALP) of each cluster is then computed by the **ALP Estimation component** using the history of all attempted goals and the corresponding outcomes. 3) In the **Prior Construction component** a cluster is sampled using the ALP estimates. The goal sampling prior is then constructed as the masking distribution assigning uniform probability over goals inside the sampled cluster (over all the goals in the history of encountered states that the clustering function would assign to this cluster) and 0 probability to goals outside the cluster. 4) A goal is then sampled, inside **The Underlying IMGEP**, from the distribution formed by combining the goal prior and the underlying IMGEP’s novelty-based goal sampling distribution i.e. a novel looking goal is sampled from the sampled cluster. For further details see section 5 and the appendix B.

The GRIMGEP framework consists of the following four components as shown in Figure 1:

I. Clustering component. It creates a clustering function that maps any goal to a cluster. Every time when a goal is to be sampled all the encountered states are passed through the backbone of a pretrained YOLO-V3[Redmon and Farhadi, 2018] object detector (taken from Guo et al. [2020], Apache License), followed by a global average pooling layer[Lin et al., 2013]. A PCA dimensionality reduction[F.R.S., 1901], and a GMM that follows it are both fit on those state embeddings. We leverage a pretrained object detector as it provides a latent space suitable for clustering different realistic objects, however this constrains our approach to realistic 3D environments. The clustering component is depicted in figure 5 of the appendix.

II. ALP Estimation component. It assigns an absolute learning progress estimate to each cluster. To estimate ALP a memory is kept containing all the goals that were sampled, the last state that was achieved while aiming for that goal, and the epoch in which this goal was sampled.

At the beginning of every epoch, the performances (rewards) corresponding to each of the previously sampled goals and the achieved last states are recomputed. We do so using the reward function from the underlying IMGEP.

By arranging the goals from the (goal, last_state) pairs history into clusters (using the clustering function fit in the clustering component) we construct a history of performances for each cluster. The performance estimates from the same epoch in the same cluster are averaged and only the last l epochs are kept in each history. The ALP is then estimated as the distance between the mean of the first and the last halves of the history as in the following equation:

$$ALP = \left| \frac{1}{l/2} \sum_{i=0}^{l/2} h_i - \frac{1}{l/2} \sum_{i=l/2}^l h_i \right| \quad (2)$$

, where l is the length of the cluster’s history, h is the cluster’s history. This component is further explained in algorithm 3 in the appendix.

III. Prior Construction component. It creates the prior distribution. A cluster is sampled according to the per cluster ALP estimates. This sampling is achieved by a Multi-Armed Bandit with ALP estimates as utilities. The goal sampling prior is then defined as a masking distribution i.e. uniform over all goals from the encountered states history that the clustering function would assign to the sampled cluster, and zero elsewhere. This distribution is the means by which the underlying IMGEP is constrained to sample a goal only from the sampled cluster.

IV. Underlying IMGEP. It can be any goal exploration process that constructs a distribution for sampling goals. The underlying IMGEPs we study (Skewfit and CountBased) both sample goals from a set of encountered image states. When used by themselves (outside the GRIMGEP framework) this

sampling is performed according to a novelty maximizing distribution p_{imggp} . However, when used inside the GRIMGEP framework, the goal sampling distribution p_{imggp} is multiplied with the prior distribution to define the final goal sampling distribution p as in the following equation:

$$p(g) = \frac{p_{prior}(g)p_{imggp}(g)}{\sum_{g \in R} p_{prior}(g)p_{imggp}(g)}, \quad (3)$$

where R is the set of all encountered states. This multiplication has a masking effect on the novelty prioritizing distribution p_{imggp} . Distribution p has non-zero probabilities only inside the cluster sampled by the Prior Construction component where they are proportional to those of p_{imggp} .

Using p instead of p_{imggp} for sampling a goal has the effect of using the underlying IMGEP only in the cluster sampled by the Prior Construction component.

6 Experiments

In this section, we present the experiments demonstrating the effectiveness of the GRIMGEP framework. We consider two novelty-seeking goal exploration processes: Skewfit[Pong et al., 2020] and CountBased (inspired by GO-Explore[Ecoffet et al., 2021]). We study their behavior when used alone and when wrapped inside the GRIMGEP framework on a 3D first-person environment containing a distractor. Additional experiments on a simpler 2D environment are explained in the appendix E.

¹

More precisely, we aim to address the following questions:

- How do current approaches behave in the presence of action-induced distractors?
- How does the GRIMGEP framework change the behavior of current approaches in the presence of noisy distractors?
- How important are the ALP estimates for the performance of the GRIMGEP framework?

6.1 Explore3D environment

Explore3D is a 3D first-person environment implemented using the MiniWorld [Chevalier-Boisvert, 2018] environment simulator (released under the Apache license). The environment consists of 14 different objects located in three separate rooms (TV room, Gallery room, Office room). The topology of the environment and the different rooms are depicted in Fig. 2. The agent’s observation space includes RGB images depicting the first-person view of the environment as shown in figures 2b-2d. The agent can interact with the environment through 7 continuous actions: four for movement, two for rotating the camera and one for turning the TV on or off. The starting position is in the TV room facing the clock as depicted by the red triangle in Fig. 8a.

At the beginning of each episode the TV is in the OFF position (a black screen). When the agent turns on the TV a random image from the ImageNet [Krizhevsky et al., 2012] validation set is sampled and shown on the screen. The turned on TV then changes the image every step with a probability of 0.1.

The environment provides a challenge for exploration in the following two aspects: 1) if no explicit exploration incentive is used the agent will not be able to explore farther from the starting TV room, and 2) novelty based exploration will be attracted to the diversity of images on the TV. Furthermore, the reaching of goals depicting a particular image on the TV is almost impossible as the agent has no control over which image, from the 50000 ImageNet validation images, will be shown on the screen.

For evaluation, we construct a test set consisting of 20 goal images depicting all the objects from the environment except the TV. The metric used for evaluation is the f1 score between the objects visible in the last state and in the goal. For further details about this metric, refer to appendix C.

6.2 Results

In the following section we present the experiments with Skewfit and CountBased. We refer to the GRIMGEP wrapped versions of those algorithms as GRIM-Skewfit and GRIM-CountBased.

¹We reused the following codebases in our experiments: <https://github.com/vitchyr/rlkit> (MIT licence), GluonCV [Guo et al., 2020] (Apache License), sklearn [Pedregosa et al., 2011] (BSD licence)




Figure 2: Different rooms of the 3D environment

How do current novelty-based approaches behave in the presence of noisy distractors?

Figure 3 shows the performance of CountBased and Skewfit, along with the corresponding percentage of all sampled goals depicting the active TV. Both baselines oversample goals depicting the active TV, thereby limiting their performances. Furthermore, as they continue to oversample those goals their performance diminishes resulting in *catastrophic forgetting*.

These experiments show that the presence of a noisy distractor is challenging for the novelty-based exploration of both Skewfit and CountBased. This motivates the use of more advanced exploration techniques such as GRIMGEP.

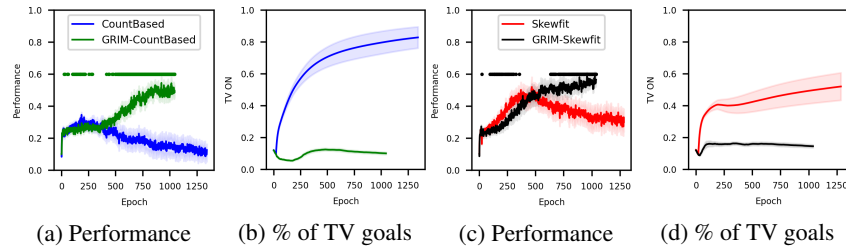


Figure 3: The performance and the percentage of sampled goals depicting the active TV (\pm std.err.). The dots depict statistical significance (Welch t-test, $p=0.05$, 15 seeds)

How does the GRIMGEP framework change the behavior of current approaches in the presence of noisy distractors?

As Skewfit tends to behave unstably if exploring only with the novelty bonuses, the authors introduced a regularization hyperparameter α . α enables interpolation between a uniform ($\alpha = 0$) and a completely exploratory distribution ($\alpha = -1$) making Skewfit more robust.

Table 1 shows the performances at epoch 1000, and the percentage of all goals sampled up to epoch 1000 that depict the active TV. This is shown for CountBased, Skewfit with two different values α , and for those three approaches wrapped using the GRIMGEP framework. We can see that, for all three approaches, GRIMGEP improves the performance by reducing the sampling of TV goals.

In figure 3 the performance and the sampling of goals depicting the active TV are shown. More precisely, in figures 3c and 3d Skewfit with $\alpha = -0.25$ is compared to GRIM-Skewfit with $\alpha = -0.75$. This choice is made to compare the best (in terms of peak performance) Skewfit with the best GRIM-Skewfit. The remaining training curves are shown in the figure 6 of the appendix.

In figure 3 we can see that GRIM-CountBased and GRIM-Skewfit do not suffer from *catastrophic forgetting* and are able to achieve greater performances compared to Skewfit and CountBased. This is explained by observing that, unlike their unwrapped versions, neither GRIM-Skewfit nor GRIM-CountBased are drawn to the distractor.

These experiments show that wrapping the novelty seeking approaches with the GRIMGEP framework alleviates the problems introduced by noisy distractors.

How important are the ALP estimates for the performance of the GRIMGEP framework?

We study this question by doing an ablation study on the GRIMGEP wrapped CountBased. Figure 4 shows how GRIM-CountBased behaves when the clusters are sampled uniformly (GRIM-UNI-CountBased) as compared to when using learning progress (GRIM-LP-CountBased). The figure shows performance and percentages off all goals sampled for the distractor(active TV), the object

	performance	% of TV ON goals
CountBased	0.147 ± 0.235	0.796 ± 0.275
GRIM-CountBased	0.477 ± 0.111	0.101 ± 0.028
Skewfit ($\alpha = -0.25$)	0.332 ± 0.245	0.45 ± 0.318
GRIM-Skewfit ($\alpha = -0.25$)	0.516 ± 0.08	0.101 ± 0.032
Skewfit ($\alpha = -0.75$)	0.369 ± 0.161	0.232 ± 0.079
GRIM-Skewfit ($\alpha = -0.75$)	0.537 ± 0.09	0.146 ± 0.036

Table 1: Performance and the percentage of all sampled goals depicting the activated TV evaluated at epoch 1000. 15 seeds were used and all the GRIM versions are statically significantly different (Welch’s t-test, $p=0.05$) from their unwrapped counterparts

visible from the starting position (clock), an object from the Gallery room (Impression), and an object from the office room (office chair), the figures for the rest of the objects are shown in the appendix 7. We can see that even when sampling clusters uniformly (GRIM-UNI-CountBased) the TV is not being oversampled. Since the agent’s starting position is that of looking at the clock, at the beginning a lot of clusters consist mostly of goals depicting the clock. This results in both low sampling of the goals with the active TV, and high sampling of the goals with the clock. By looking at the sampling of goals with Impression and the office desk, we can see that GRIM-UNI-Countbased is not able to explore further away from the clock as opposed to GRIM-LP-CountBased.

These experiments show that, for CountBased, LP plays a significant role in the final performance of GRIMGEP by pushing the exploration farther from the starting position while at the same time not sampling the TV.




Figure 4: The performance and the percentage of sampled goals depicting the various objects (\pm std. err.) The dots depict statistical significance (Welch t-test, $p=0.05$, 15 seeds)

7 Conclusion

We studied the problem of unsupervised image-based goal exploration in first-person 3D environments with action-induced noisy distractors. Having demonstrated that current novelty-based approaches fail in such environments, we introduced the GRIMGEP framework. GRIMGEP can be wrapped around any novelty seeking goal exploration process enabling it to ignore the distractor. It detects the noisy distracting regions of the environment by estimating absolute learning progress, and uses the novelty seeking exploration only in the interesting regions. We have shown that wrapping state-of-the-art exploration algorithms inside the GRIMGEP framework allows them to ignore the distracting regions to ultimately improve their performance and alleviate catastrophic forgetting.

Limitations The limitation of this work is that the experiments are constrained to a simulated environment. It remains an open question to see how the approach would behave in the real physical world with more sophisticated kinds of distractors.

Negative societal impact The autonomous exploration approach presented here focuses on discovering diverse behaviors and do not consider human preferences. That means that such algorithms when run in the real world would not have any bias for moral or socially acceptable behavior. It is, therefore, important that in the future such methods are coupled with mechanisms enabling alignment with human values and preferences.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes] refer to paragraph "Limitations" in the conclusion
 - (c) Did you discuss any potential negative societal impacts of your work? - understand more autonomous learning agents if they are put in real world open environment [Yes] refer to paragraph "Negative societal impact" in the conclusion
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [N/A]
 - (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] As an URL in the abstract
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] Refer to the appendix D
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] Refer to figure captions.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes] Pong et al. [2020], Chevalier-Boisvert [2018], Krizhevsky et al. [2012] and Guo et al. [2020] are all cited in the paper.
 - (b) Did you mention the license of the assets? [Yes] Refer to the footnote on page 1 for Pong et al. [2020], and section 6.1 for Chevalier-Boisvert [2018], section 5 under "clustering component" for Guo et al. [2020]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] The code. Containing a new environment. Refer to the abstract for the URL.
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [No] As the only data we used was ImageNet [Krizhevsky et al., 2012], we didn't find it nessecary.

References

- Dmitry Kalashnikov, Alex Irpan, Peter Pastor, Julian Ibarz, Alexander Herzog, Eric Jang, Deirdre Quillen, Ethan Holly, Mrinal Kalakrishnan, Vincent Vanhoucke, et al. Qt-opt: Scalable deep reinforcement learning for vision-based robotic manipulation. *arXiv preprint arXiv:1806.10293*, 2018.
- Tuomas Haarnoja, Sehoon Ha, Aurick Zhou, Jie Tan, George Tucker, and Sergey Levine. Learning to walk via deep reinforcement learning. *arXiv preprint arXiv:1812.11103*, 2018a.
- Yuxiang Yang, Ken Caluwaerts, Atil Iscen, Tingnan Zhang, Jie Tan, and Vikas Sindhwani. Data efficient reinforcement learning for legged robots. In *Conference on Robot Learning*, pages 1–10, 2020.
- Shixiang Gu, Ethan Holly, Timothy Lillicrap, and Sergey Levine. Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In *2017 IEEE international conference on robotics and automation (ICRA)*, pages 3389–3396. IEEE, 2017.
- 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.
- Vitchyr H. Pong, Murtaza Dalal, Steven Lin, Ashvin Nair, Shikhar Bahl, and Sergey Levine. Skew-fit: State-covering self-supervised reinforcement learning, 2020.
- Soroush Nasiriany, Vitchyr H Pong, Ashvin Nair, Alexander Khazatsky, Glen Berseth, and Sergey Levine. Disco rl: Distribution-conditioned reinforcement learning for general-purpose policies. *arXiv preprint arXiv:2104.11707*, 2021.
- Adrien Laversanne-Finot, Alexandre Péré, and Pierre-Yves Oudeyer. Curiosity driven exploration of learned disentangled goal spaces. *arXiv preprint arXiv:1807.01521*, 2018.
- Cédric Colas, Pierre Fournier, Mohamed Chetouani, Olivier Sigaud, and Pierre-Yves Oudeyer. Curious: intrinsically motivated modular multi-goal reinforcement learning. In *International conference on machine learning*, pages 1331–1340, 2019.
- Adrien Ecoffet, Joost Huizinga, Joel Lehman, Kenneth O. Stanley, and Jeff Clune. First return, then explore. *Nature*, 590(7847):580–586, Feb 2021. ISSN 1476-4687. doi: 10.1038/s41586-020-03157-9. URL <https://doi.org/10.1038/s41586-020-03157-9>.
- Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, Pieter Abbeel, and Wojciech Zaremba. Hindsight experience replay. *CoRR*, abs/1707.01495, 2017. URL <http://arxiv.org/abs/1707.01495>.
- Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, and Sergey Levine. Soft actor-critic algorithms and applications. *CoRR*, abs/1812.05905, 2018b. URL <http://arxiv.org/abs/1812.05905>.
- David Warde-Farley, Tom Van de Wiele, Tejas D. Kulkarni, Catalin Ionescu, Steven Hansen, and Volodymyr Mnih. Unsupervised control through non-parametric discriminative rewards. *CoRR*, abs/1811.11359, 2018. URL <http://arxiv.org/abs/1811.11359>.
- Frédéric Kaplan and Pierre-Yves Oudeyer. The progress-drive hypothesis: an interpretation of early imitation. In K. Dautenhahn and C. Nehaniv, editors, *Models and mechanisms of imitation and social learning: Behavioural, social and communication dimensions*. Cambridge University Press, 2007. Language.
- Yuri Burda, Harri Edwards, Deepak Pathak, Amos Storkey, Trevor Darrell, and Alexei A. Efros. Large-scale study of curiosity-driven learning. In *ICLR*, 2019.
- Jürgen Schmidhuber. Curious model-building control systems. In *Proc. international joint conference on neural networks*, pages 1458–1463, 1991.

- Pierre-Yves Oudeyer, Frederic Kaplan, and Verena V Hafner. Intrinsic motivation systems for autonomous mental development. *IEEE transactions on evolutionary computation*, 11(2):265–286, 2007.
- 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. ISSN 09218890. doi: 10.1016/j.robot.2012.05.008.
- Sébastien Forestier, Yoan Mollard, and Pierre-Yves Oudeyer. Intrinsically motivated goal exploration processes with automatic curriculum learning. *CoRR*, abs/1708.02190, 2017. URL <http://arxiv.org/abs/1708.02190>.
- Matthias Rolf, Jochen J Steil, and Michael Gienger. Goal babbling permits direct learning of inverse kinematics. *IEEE Transactions on Autonomous Mental Development*, 2(3):216–229, 2010.
- Sao May Nguyen and Pierre-Yves Oudeyer. Socially guided intrinsic motivation for robot learning of motor skills. *Autonomous Robots*, 36(3):273–294, 2014.
- Joel Lehman and Kenneth O Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, 19(2):189–223, 2011.
- Benjamin Eysenbach, Abhishek Gupta, Julian Ibarz, and Sergey Levine. Diversity is all you need: Learning skills without a reward function. *arXiv preprint arXiv:1802.06070*, 2018.
- Marc Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Remi Munos. Unifying count-based exploration and intrinsic motivation. In *Advances in neural information processing systems*, pages 1471–1479, 2016.
- Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 16–17, 2017.
- Yuri Burda, Harrison Edwards, Deepak Pathak, Amos J. Storkey, Trevor Darrell, and Alexei A. Efros. Large-scale study of curiosity-driven learning. *CoRR*, abs/1808.04355, 2018a. URL <http://arxiv.org/abs/1808.04355>.
- Roberta Raileanu and Tim Rocktäschel. RIDE: rewarding impact-driven exploration for procedurally-generated environments. *CoRR*, abs/2002.12292, 2020. URL <https://arxiv.org/abs/2002.12292>.
- Yuri Burda, Harrison Edwards, Amos Storkey, and Oleg Klimov. Exploration by random network distillation. *arXiv preprint arXiv:1810.12894*, 2018b.
- Kuno Kim, Megumi Sano, Julian De Freitas, Nick Haber, and Daniel Yamins. Active world model learning with progress curiosity. *arXiv preprint arXiv:2007.07853*, 2020.
- Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *CoRR*, abs/1804.02767, 2018. URL <http://arxiv.org/abs/1804.02767>.
- Jian Guo, He He, Tong He, Leonard Lausen, Mu Li, Haibin Lin, Xingjian Shi, Chenguang Wang, Junyuan Xie, Sheng Zha, Aston Zhang, Hang Zhang, Zhi Zhang, Zhongyue Zhang, Shuai Zheng, and Yi Zhu. Gluoncv and gluonnlp: Deep learning in computer vision and natural language processing. *Journal of Machine Learning Research*, 21(23):1–7, 2020. URL <http://jmlr.org/papers/v21/19-429.html>.
- Min Lin, Qiang Chen, and Shuicheng Yan. Network in network. *arXiv preprint arXiv:1312.4400*, 2013.
- Karl Pearson F.R.S. Liii. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):559–572, 1901. doi: 10.1080/14786440109462720.

- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Maxime Chevalier-Boisvert. gym-miniworld environment for openai gym. <https://github.com/maximecb/gym-miniworld>, 2018.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012.
- Tsung-Yi Lin, Michael Maire, Serge J. Belongie, Lubomir D. Bourdev, Ross B. Girshick, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312, 2014. URL <http://arxiv.org/abs/1405.0312>.
- Carl Edward Rasmussen. The infinite gaussian mixture model. In *Advances in neural information processing systems*, pages 554–560, 2000.
- Hamparsum Bozdogan. Model selection and akaike’s information criterion (aic): The general theory and its analytical extensions. *Psychometrika*, 52:345–370, 02 1987. doi: 10.1007/BF02294361.
- Rémy Portelas, Cédric Colas, Katja Hofmann, and Pierre-Yves Oudeyer. Teacher algorithms for curriculum learning of deep rl in continuously parameterized environments. In Leslie Pack Kaelbling, Danica Kragic, and Komei Sugiura, editors, *Proceedings of the Conference on Robot Learning*, volume 100 of *Proceedings of Machine Learning Research*, pages 835–853. PMLR, 30 Oct–01 Nov 2020. URL <http://proceedings.mlr.press/v100/portelas20a.html>.

A Additional details on the studied underlying IMGEPs

A.1 Skewfit algorithm

Skew-fit was proposed in [Pong et al., 2020] primarily as an iterative algorithm that aims to train a generative model modeling the uniform distribution over the feasible goal space. To ensure this feasibility, Skew-fit constructs a dataset for training the generative model only from the already observed states. They can do this because they, like us, assume the goal space to be equivalent to the state space. After each epoch of the generative model training, a new Skewed dataset is constructed. It is constructed by setting each state’s weight (probability of being sampled) as inversely proportional to the current generative model’s probability of that state. This skewed dataset is then used to train the generative model in the next epoch. This dataset’s entropy increases in each epoch, so the model converges to modeling a uniform distribution.

This algorithm is then applied to unsupervised goal-driven exploration. The generative model is a VAE, and for its training dataset the replay buffer is used. They propose two mechanisms for goal sampling: sampling directly from the current VAE (which they call sampling from the generative model q_{ψ}^G), and sampling from the skewed replay buffer (which they call sampling from the skewed distribution p_{skewed}). In this work, though both mechanisms could be used with the GRIMGEP framework, we focus on the latter as we found it worked better. This is aligned with the conclusion from [Pong et al., 2020] as they use p_{skewed} for more complex environments.

They train a goal-conditioned policy using SAC[Haarnoja et al., 2018b] coupled with HER [Andrychowicz et al., 2017] with the reward being the current VAE’s negative L2 latent distance between the goal and the state.

In every epoch, goals are sampled, episodes run, new data added to the replay buffer, and both the VAE and the agent are trained. It should be noted that for training the SAC agent, data is sampled uniformly from the replay buffer, but for training the VAE and sampling of replacement goals for HER, the sample sampling mechanism as for goal sampling is used (a skewed replay buffer).

They denote this algorithm Skewfit + RIG though for simplicity we denote it Skewfit.

About the α hyperparameter It is relevant to note that Skewfit has a regularization hyperparameter α which interpolates between an uniform and a skewed distribution (-1 being completely skewed and 0 being uniform). This is done by exponentiating the original probability with α and then normalizing. In their work they experiment with 4 different values of this parameter (-0.25, -0.5, -0.75, -1.0)[Pong et al., 2020].

In our experiments we found that -0.25 achieves better peak performance for Skewfit, but when wrapped inside the GRIMGEP frames -0.75 works better. As can be seen in figure 6.

That demonstrates another benefit of the GRIMGEP framework. Using Skewfit inside the GRIMGEP framework enables us to use more aggressive exploration bonuses ($\alpha = -0.75$) without training becoming unstable.

A.2 CountBased approach

This approach is very similar to Skewfit, but it prioritizes goal sampling using a different intrinsic reward. This intrinsic reward is computed using a count-based novelty measure. As shown in equation 4, we downsize an image (to 3x3) and quantize each channel into 4 different values (4^3 different colors), then we count how many times the representation was observed. The goal’s intrinsic reward is then computed by exponentiating this count by a hyperparameter $\alpha \in [-1, 0]$, as in [Pong et al., 2020]. As in Skewfit, we use this reward to construct the distribution to prioritize data for training the VAE, sampling exploration goals and sampling HER replacement goals. When using this approach as part of the GRIMGEP framework we multiply this distribution with the prior as in equation 3.

$$R_{CB}(image) = count(quantize(downsize(image)))^\alpha \tag{4}$$

B Details on the GRIMGEP approach

In this section we will present additional details of the GRIMGEP framework.

The biased distribution, p which was constructed by Prior construction component, is used everywhere where the underlying IMGEP would normally use p_{imgep} . Therefore, for three purposes: 1) sampling goals for exploration, 2) sampling replacement goals for HER, and 3) since all of our underlying IMGEPs train their own VAE, biasing the training of this VAE.

A more detailed description of shown in the algorithm 2. First, we do a warmup phase to fill the replay buffer by doing random actions. Then, for the first *start_exploration* epochs, we do not use any kind of intrinsic rewards but just resample goals uniformly from the replay buffer. After this, we start using intrinsic rewards to prioritize the goal sampling. In practice, we make use of parallel computing by sampling ten goals per epoch. For each of those goals, we sample a new cluster, construct a new prior and then sample a new goal using the probability distribution obtained by combining the prior and the goal sampling probability distribution of the IMGEP. This is depicted by the algorithm 2 where the blue colors depict GRIMGEP addition.

Algorithm 2: Intrinsically Motivated Goal Exploration Process in the Deep RL setting (blue lines correspond to GRIMGEP additions)

```

π.init()  $\mathcal{R}$ .init()           // goal-conditioned policy and the reward function (VAE)
encountered_states = []
performance_history = []      // A history of (sampled_goal, last_state) pairs
// random acting rollouts to fill the replay buffer
for  $N_{warmup}$  do
    traj = env.random_rollout()
    encountered_states.add_states(traj)
// exploration
for ep in  $N_{epochs}$  do
    for  $N_{parallel\_processes}$  do
        if ep > start_exploration then
            // novelty-based exploration of the underlying IMGEP
            goal_intrinsic_rewards = compute_intrinsic_rewards(encountered_states)

            // Assign a cluster_id to every goal from the performance_history
            cluster_ids = clustering_fn(performance_history)

            // Recompute the performances and estimate per-cluster ALP
            per_cluster_alp = estimate_ALPs(cluster_ids, performance_history,  $\mathcal{R}$ ) // by algo. 3
            sampled_cluster = sample(clusters, per_cluster_alp) // by eq. 5

            // Create a masking distribution by eq. 6
            prior = construct_prior(sampled_cluster, encountered_states)

            // Integrate the bias by eq. 3
            goal_intrinsic_rewards *= prior
        else
            goal_intrinsic_rewards = [1,1...1] // at the beginning sample uniformly
        goal = prioritized_sample_goal(encountered_states, goal_intrinsic_rewards)

        trajectory = env.policy_rollout(π, goal)
        encountered_states.add_states(trajectory)
        performance_history.add_pair((goal, trajectory.last_state))
     $\mathcal{R}$ .train() π.train()           // fit reward function (VAE) and policy
    clustering_fn.train(encountered_states) // fit PCA and GMM

```

- I. **Clustering component:** The pipeline of this component is shown in figure 5. The pretrained YOLO-v3 object detector whose backbone was used was taken from the GluonCV[Guo et al., 2020] library ('yolo3_darknet53_coco'). This model was pretrained on the COCO[Lin et al., 2014] object detection dataset (Creative Commons 4.0).
- II. **ALP Estimation component:** This component's pseudocode is shown in the algorithm 3. The whole history of (goal, last_state) pairs is separated by the cluster in which the goal was assigned to. The performance estimates are recomputed as the reward function is the $L2$




Figure 5: The clustering component. A goal image is passed through the backbone of a pretrained yolo detector followed by a global average pooling operation. The dimensionality of this vector is then reduced using a PCA transformation and a GMM is used to assign the goal to a cluster.

distance in the latent space of a VAE trained online. The performances corresponding to goals sampled in the same epoch and assigned to the same cluster are averaged to create the per-cluster histories. Only the estimates from the last l epochs are kept. These histories then used to estimate absolute learning progress for each cluster by equation 2.

Algorithm 3: ALP estimation component

```

def estimate_ALP(cluster_ids, performance_history,  $\mathcal{R}$ ):
    // Initialize the per-cluster histories as an empty list.
     $h_c = [], \forall c \in \{1..n\_clusters\}$ 

    // Performance computation
    // Estimate performances epoch by epoch
    for epoch  $\in$  epochs_so_far do
         $p_c = [], \forall c \in \{1..n\_clusters\}$ ; // per-cluster performances are reset
        // Extract the (goal, last_state) pairs collected in this epoch
        epoch_rollouts = performance_history[epoch]
        // compute the performances from this epoch for each cluster
        for rollout  $\in$  epoch_rollouts do
            c = cluster_ids[rollout.goal]
            goal_perf =  $\mathcal{R}$ (rollout.goal, rollout.last_state); // recompute the rewards
             $p_c.append(goal\_perf)$ 
        // add the mean performances from this epoch to the clusters' history
        for c  $\in$  n_clusters do
             $h_c.append(mean(p_c))$ 

    // LP estimation
    for c  $\in$  clusters do
         $ALP_c = estimate\_ALP(h_c)$ ; // by equation 2

```

III. Prior Construction component:

Our prior takes the form of a masking distribution. It is constructed in the following steps:

1. Using per-cluster ALP estimates, sample a cluster according to the probability defined in the following equation:

$$p(c) = \frac{4}{5} \frac{LP_c^T}{\sum_{i=1}^C LP_c^T} + \frac{1}{5} \frac{1}{C} \quad (5)$$

,where C is the number of clusters and T is a hyperparameter. This is a Multi-Armed Bandit with a 20% uniform prior. The hyperparameter T enables us to set how much priority we want to give to the high ALP clusters, which is useful if we have a lot of clusters with small LPs.

2. Construct the prior according to the following equation:

$$prior(g) = \begin{cases} \frac{1}{n_c}, & cl(g) == c \\ 0, & else \end{cases} \quad (6)$$

, where c is the sampled cluster, n_c is the number of goals from the replay buffer that would be assigned to cluster c , and cl is the clustering function assigning any state to a cluster. This is essentially a masking distribution giving uniform probabilities for goals inside the sampled cluster and zero for the rest. Another way of looking at this prior is saying what we sample a cluster and allow the underlying IMGEP to sample the goals only from this cluster.

B.1 Additional details about the experiments

In this section, we will show some additional figures explaining the results discussed in the main text 6.

Figure 6 shows the success rates and the percentage of all sampled goals that depict the active TV for CountBased, Skewfit with two different values of hyperparameter α , compared to those three approaches wrapped inside the GRIMGEP framework. We can see that GRIMGEP improves the performance for all three underlying IMGEPs and doesn't suffer from catastrophic forgetting. We can also see that, even though Skewfit with $\alpha = -0.75$ eventually stops oversampling goals with the active TV, its performance is still diminished as a consequence of oversampling in the beginning.

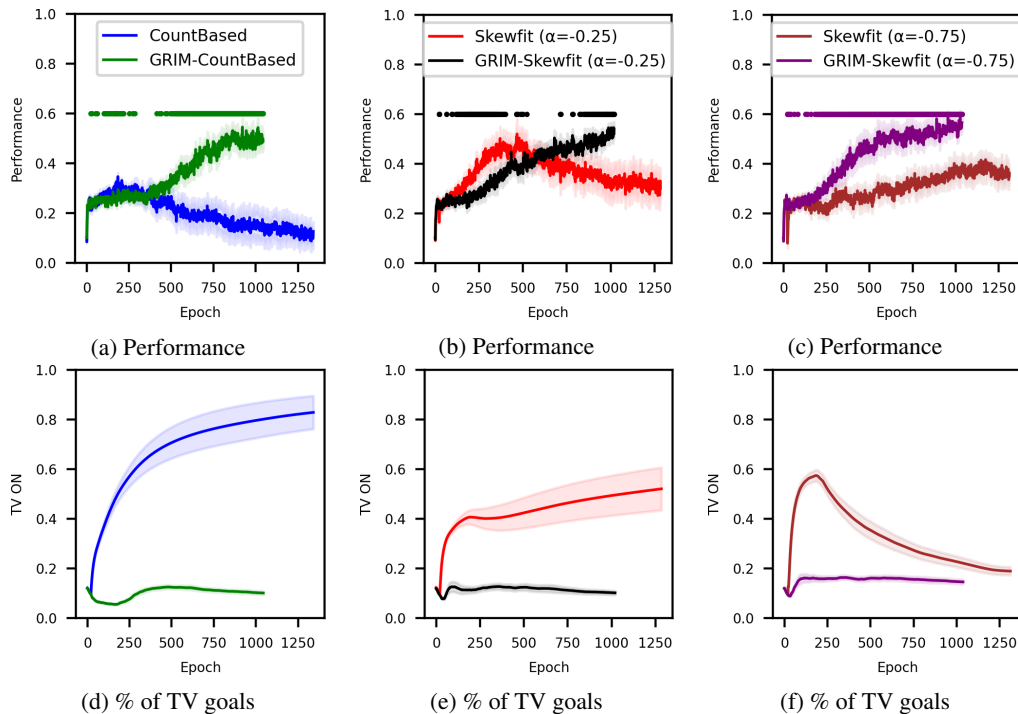


Figure 6: The performance and the percentage of sampled goals depicting the active TV (\pm std. err.). Shown for Countbased, Skewfit with two different values of α , compared to the three approaches wrapped inside the GRIMGEP framework. We can see that the performance of all the three approaches is improved, and catastrophic forgetting avoided, when wrapped inside the GRIMGEP framework. The dots depict statistical significance (Welch t-test, $p=0.05$, 15 seeds)

Figure 7 shows the success rate and the goal sampling statistics for every object in the environment. It compares the CountBased approach to its GRIMGEP wrapped counterpart. The first row shows objects from the TV room, the second row objects from the GALLERY room, and the third row objects from the OFFICE room. We can see that if ALP is not used, i.e. the clusters are sampled uniformly, the agent continues to sample many goals from the TV room, especially goals depicting the clock. This demonstrates the inability of the agent to explore further from the starting position (the starting position is in the TV room looking at the clock). Furthermore, we can see that the objects from the other rooms are sampled more if LP is used.

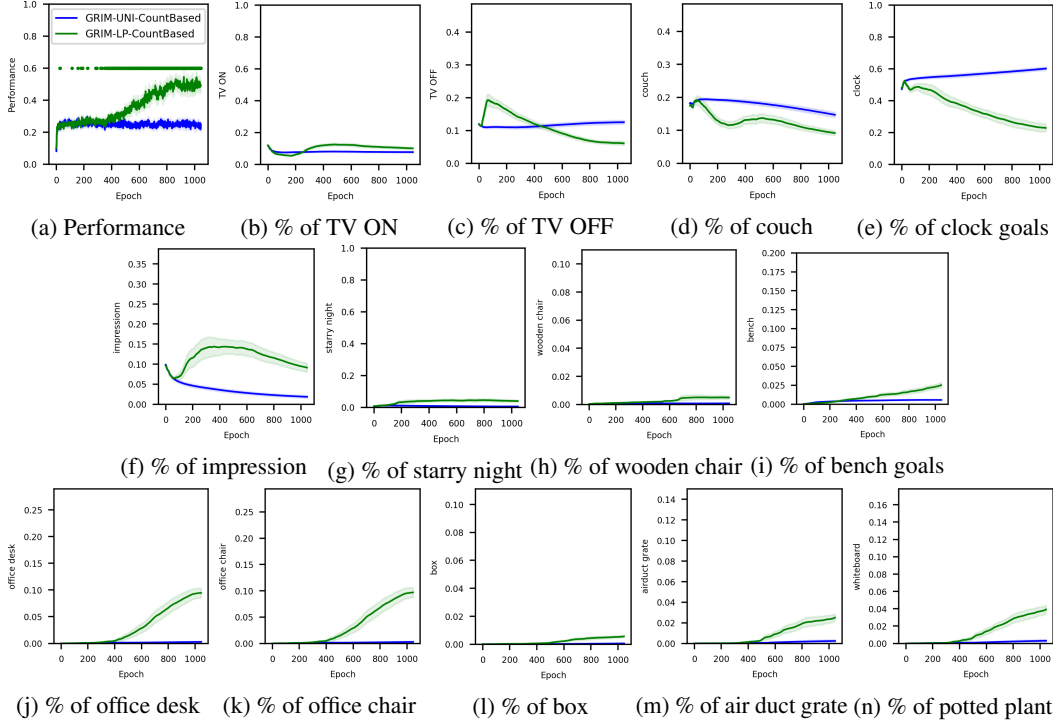


Figure 7: The Performance and the sampling distribution of all the objects in the environment. Shown for GRIMGEP wrapped CountBased, as compared to GRIMGEP wrapped CountBased which doesn't use absolute learning progress(ALP), i.e. samples the cluster uniformly. We can see that ALP enables the agent to move from the starting TV room (figures b-e) to the other rooms of the environment(f-n), thereby reaching superior performances (figure a). The dots depict statistical significance (Welch t-test, $p=0.05$, 15 seeds).

C Details on the evaluation metric

To evaluate the fulfillment of a goal, we construct the *visible entities f1 score* metric. Access to this metric is only available in evaluation. Using the environment simulator, a list of visible objects (entities) is extracted for both the goal and the last state from the episode. By treating objects visible in the goal as ground truth, precision and recall, and finally f1 score are calculated.

This metric, although it captures well the task of "seeing the goal image", is not perfect. We can notice that it is insensitive both to the angle and the distance at which the agent is looking at the objects. Furthermore, in some cases an object may be visible on the state image with only a pixel of two. This difference, even though it's very hard for the agent to fix, will be treated as if the object is fully present. Nonetheless, we found the metric sufficiently accurate for comparing different algorithms.

D Hyperparameters

For the underlying IMGEPs, all the hyperparameters, including the ones for the training VAE, are the same as in Pong et al. [2020] in the "Visual Door" experiments, except the α hyperparameter (see Appendix A.1).

Other hyperparameters are shown in table 2.

E Preliminary experiments in the 2D PlaygroundRGB environment

The preliminary experiments used in constructing the GRIMGEP framework were done on a simpler 2D environment called PlaygroundRGB. Since some changes were made in the architecture for

T	5
episode length	50
l - cluster history length	50
d - PCA latent size	25
k - GMM number of clusters	25

Table 2: Hyperparameters.

extending the agent to the 3D environment, the GRIMGEP architecture used in these experiments on PlaygroundRGB environment is not the same as the one described in the main paper.

In this section we will explain the environment those differences and the experiments.

E.1 Changes to the GRIMGEP between the experiments in the PlaygorunrRGB(2D environment) an the Explore3D

Clustering component Since the 2D environment does not have realistic objects and has a simpler underlying latent space an additional VAE was trained to create the clustering latent space on which to train the GMM i.e. the YOLO backbone, global average pooling, and the PCA were replaced with this VAE. This additional VAE is not the same VAE as the one used inside the underlying IMGEP to compute the rewards for training the policy. It is important that this VAE’s latent space creates the features which are relevant for separating the regions and not on the specific details relevant for training the agent. For this purpose, we reduce the size of the VAE and its latent space. For the hyperparameters, see table 3. We train this VAE online after each epoch on the data uniformly sampled from the replay buffer.

In the experiments on the PlaygroundRGB environment ten different GMM[Rasmussen, 2000] models each having a different number of clusters (1, 3, . . . , 19) was trained. Then we the best one was chosen by their AIC[Bozdogan, 1987] score. This mechanism was inspired by the one used in Portelas et al. [2020]. The best GMM was set as the clustering function. Each epoch the process is repeated and a new GMM selected.

In the experiments on the Explore3D environment, the number of clusters was simply treated as the hyperparameter (k) as the procedure described above introduced unnecessary computation.

The clustering VAE hyperparameterss are shown in table 3.

representation size	3
batch_size	128
beta	1
lr	0.001
Encoder	
kernel sizes	[5, 3]
num of channels	[4, 4]
strides	[3, 2]
Decoder	
kernel sizes	[3, 3]
num of channels	[4, 4]
strides	[2, 2]

Table 3: The Clustering VAE hyperparmaeters. A reduced version of the training VAE used in the underlying IMGEPs.

E.2 PlaygroundRGB

The *PlaygroundRGB* environment consists of three rooms through which the agent can move. The agent’s actions space consists of 3 continuous actions which control the position of the gripper and can close/open /the arm. It observes the environment as an image of a top down view of the current room.

The available rooms are depicted in Fig. 8 and the topology of the environment in Fig. 8a. The agent always starts in the Start room. All the possible goals inside this room are very easy and require only moving the gripper to the correct location. The Object room represents the *interesting* part of the environment as it contains a movable object and is the only non-distracting part of the environment. The TV room plays the role of an action induced noisy distractor. This room contains a TV that can be turned on by closing the gripper. When the TV is turned ON the location of the TV and the background color are randomized (a random color from a set of 5).

In short, to solve the task, exploration should guide the agent out of the Start room, away from the TV room, and into the Object room.

Since goals from the Start room are easy, we expect any goal exploration algorithm to learn goals inside this room. However, only algorithms exploring well should master goals inside the Object room. This is why, for evaluation, we construct a static test set of 25 goals from the Object room. Goals are completed if in the last state both objects are in the correct location. The performance of the agent is the average success over this evaluation set.

For evaluation, the 25 test goals from the OBJECT room were constructed by selecting 5 possible locations (center, NW, NE, SW, SE) and doing the Cartesian product of these locations for both objects. We evaluate the location of an object as correct if its L_∞ distance from the goal’s location is less than 0.2.




Figure 8: Different rooms of the PlaygroundRGB environment. A preliminary 2D environment.

E.2.1 Experiments

In these experiments we address the same questions as in the Explore3D. Those questions are:

- How do current approaches behave in the presence of action-induced distractors?
- How does the GRIMGEP framework change the behavior of current approaches in the presence of noisy distractors?
- How important are the ALP estimates for the performance of the GRIMGEP framework?

We study this questions using two novelty seeking IMGEPs (Skewfit and CountBased) and one which doesn’t have exploration incentives and samples goals uniformly from the history of encountered states (OnlineRIG).

How do current approaches behave in the presence of noisy distractors?

We can see, in Fig. 9, that Skewfit is heavily drawn to the noisy part of the TV room resulting in very low sampling of other parts of the environment, notably the Object room. As a result of not exploring the Object room enough, the final performance diminishes.

CountBased also samples a lot of goals in the TV room. However, in comparison to Skewfit, the focus is separated between both the TV-on and TV-off goals (see Fig. 9d and Fig. 9e).

OnlineRIG is also not able to achieve high performances. The reason is that, due to the lack of exploration incentive, it samples goals mostly from the Start room as can be seen in Figure 11).

Overall, this experiment demonstrates that this environment requires exploration incentives but, that these incentives should not be novelty based.

How does the GRIMGEP framework change the behavior of current approaches in the presence of noisy distractors?

To answer this question we wrap the Skewfit and CountBased with the GRIMGEP framework. As can be seen in Fig. 9, both GRIM-Skewfit and GRIM-CountBased, focus more on the Object room and less on the TV room than their unwrapped counterparts. This results in much better performances at the end of the training. When using OnlineRIG inside the GRIMGEP framework we can again observe a strong focus on the Object room (see Fig. 11b), but we can also see that this focus alone is not sufficient to greatly improve performance.

This leads us to the conclusion that GRIMGEP successfully does two things: 1) it detects the relevant part of the goal space, and 2) successfully uses the novelty seeking exploration in this relevant region.

How important are the ALP estimates for the performance of the GRIMGEP framework?

We study this question by doing an ablation study on the cluster sampling technique. In the experiments in Figure 10a we test how cluster sampling based on LP (GRIM-LP-*imgep_name*) differs in comparison to uniform cluster sampling (GRIM-UNI-*imgep_name*). As shown in Fig 10a, GRIMGEPs that use LP outperform the GRIMGEPs that don't. Furthermore, we can see that, when sampling the clusters uniformly, the Object room is sampled considerably less.

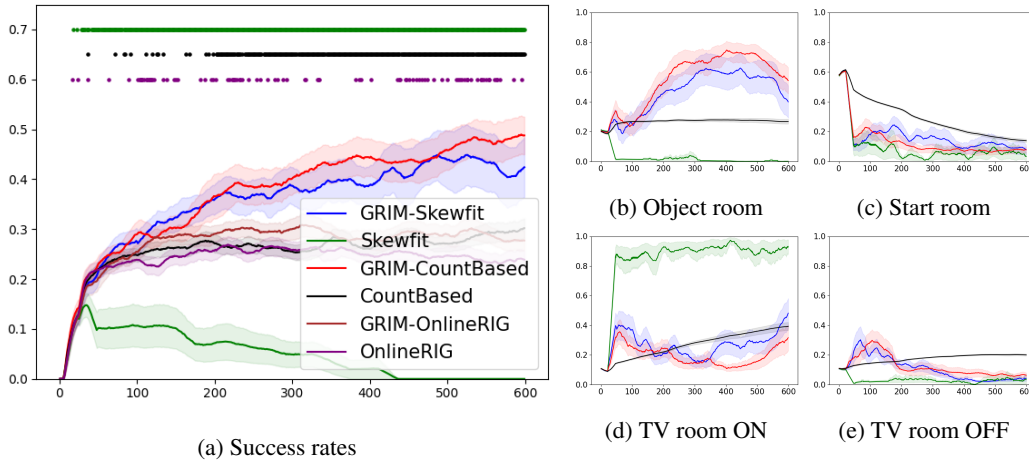


Figure 9: Comparison of Countbased Skewfit and OnlineRIG when used alone and in combination with when used inside the GRIMGEP framework on the preliminary 2D PlaygroundRGB environment. We can see that GRIMGEP improves the performance of both Skewfit, CountBased and OnlineRIG. Furthermore, we can see that the performance of GRIM-OnlineRIG is inferior to that of GRIM-Skewfit and GRIM-CountBased. Since OnlineRIG doesn't use exploration bonuses (samples goals uniformly), this shows that GRIMGEP is able to both detect the interesting region of the goal space and reuse the novelty-based exploration of CountBased and Skewfit in this region. Ten seeds were used, and the dots depict statistically significant ($p < 0.05$, Welch's t-test) results. The shaded areas correspond to standard errors and the bold line to the mean (smoothed over 25 epochs).

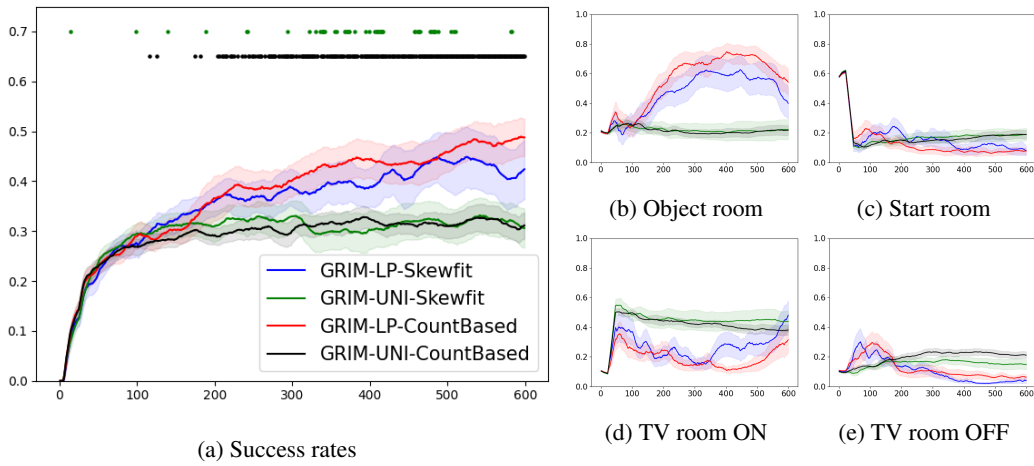


Figure 10: Comparison of the GRIMGEP framework with the sampling of clusters according to ALP and uniformly on the preliminary 2D PlaygroundRGB environment. It is visible that GRIMGEP works better when LP is used to select the most interesting cluster. The dots depict statistically significant ($p < 0.05$, Welch’s t-test) results when compared to the LP version (GRIM-LP-*imgep_name*). The shaded areas correspond to standard errors and the bold line to the mean (smoothed over 25 epochs).

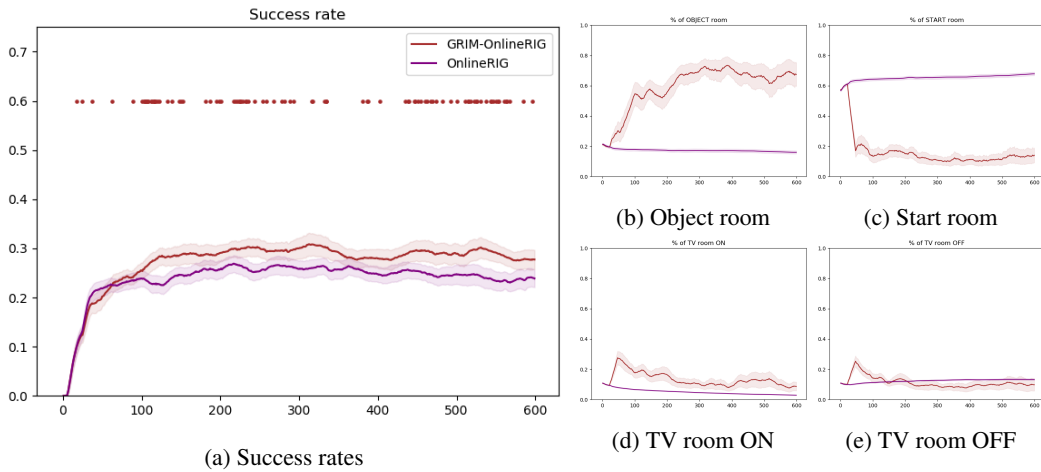


Figure 11: Comparison of Online-RIG alone and in combination with the GRIMGEP framework. On the preliminary 2D PlaygroundRGB environment. We can see that GRIMGEP improves the performance of OnlineRIG by focusing on the Object room. Ten seeds were used, and the dots depict statistically significant ($p < 0.05$, Welch’s t-test) results. The shaded areas correspond to standard errors and the bold line to the mean (smoothed over 25 epochs).