# Exploring and Reaching Goals via World Models

**Anonymous Author(s)**
Affiliation
Address
email

## Abstract

How can an artificial agent learn to solve a wide range of tasks in a complex visual environment in the absence of external supervision? We decompose this question into two problems, global exploration of the environment and learning to reliably reach situations found during exploration. We introduce the Explore Achieve Network (ExaNet), a unified solution to these by learning a world model from the high-dimensional images and using it to train an explorer and an achiever policy from imagined trajectories. Unlike prior methods that explore by reaching previously visited states, our explorer plans to discover unseen surprising states through foresight, which are then used as diverse targets for the achiever. After the unsupervised phase, ExaNet solves tasks specified by goal images without any additional learning. We introduce a challenging benchmark spanning across four standard robotic manipulation and locomotion domains with a total of over 40 test tasks. Our agent substantially outperforms previous approaches to unsupervised goal reaching and achieves goals that require interacting with multiple objects in sequence. Finally, to demonstrate the scalability and generality of our approach, we train a single general agent across four distinct environments. For videos, see https://sites.google.com/view/exanet/home.

## 1 Introduction

This paper tackles the question of how to build an autonomous agent that can achieve a diverse set of tasks specified by a user at test time, such as a legged robot standing in certain pose, a robotic arm rearranging blocks between bins, or performing chores in a kitchen. Such a general system would be difficult to realize within the traditional reinforcement learning (RL) paradigm, where a user specifies task rewards that the agent finds by sampling its environment. Designing task rewards requires domain knowledge, is time-consuming, and prone to human errors. Moreover, a traditional RL would have to explore and retrain for every new task. Instead, the problem of solving diverse tasks is often studied under the paradigm of unsupervised RL, where an agent builds a repertoire of skills without any supervision, which later enable solving human-specified goals with no or limited further training.

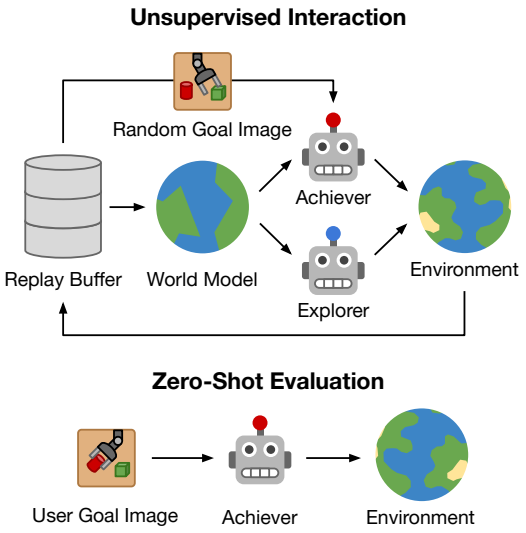


Figure 1: ExaNet learns a world model without any supervision, and leverages it to train two policies in imagination. The *explorer* finds new images and the *achiever* learns to reliably reach them. Once trained, the achiever reaches user-specified goals zero-shot without further training at test time.

**Two challenges** Building a capable unsupervised agent for reaching arbitrary goals presents two major challenges. First, the user-specified

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

goals are often diverse and rare situations, hence the agent needs to globally explore its environment. Second, the agent then needs to learn to reliably achieve the diverse intrinsic goals it found during exploration to prepare itself for user-specified goals at test time. We propose a unified approach that learns a world model together with an explorer policy and a goal achiever policy in a fully unsupervised manner to address both challenges.

**Goal exploration** Prior methods for unsupervised goal reaching approach the exploration problem by relabeling trajectories with previously visited states as goals [2, 48] or by sampling goals from a density model of previous inputs [31]. The goals can be sampled uniformly [31], or rarely visited goals can be oversampled [13, 36, 52]. These approaches have in common that they explore by visiting goals that either have been reached before or are interpolations of previously reached states. They do not explore far beyond the frontier and suffer from a chicken and egg problem: the policy does not yet know how to reach interesting goals and thus repeats already known behaviors. We observe that this leads to poor exploration performance in such methods. To rectify this issue, we leverage a learned world model to train an *explorer* policy in imagination. The world model enables the explorer to target novel states that are several steps away from the frontier [41, 42], unlocking successful goal reaching in complex environments where exploration is nontrivial.

**Goal reaching** The diverse experience collected by the explorer provides start and goal states for training the goal *achiever* policy. Because our approach does not rely on goal relabeling [2], we are free to train the achiever using on-policy trajectories generated by the world model [23, 24, 27, 45]. Training the achiever requires measuring the distance between an the states along an imagined trajectory and the goal, a non-trivial problem when inputs are high-dimensional images. We empirically analyze two choices for the distance measures, namely the cosine distance in latent space and a learned temporal distance and provide recommendations on which distance function is appropriate for different settings.

**Contributions** The introduce the Explore Achieve Network (ExaNet), an unsupervised goal reaching agent that trains an *explorer* and an *achiever* within a shared world model. At test time, the achiever solves challenging locomotion and manipulation tasks provided as user-specified goal images. For a thorough evaluation, we introduce a new challenging goal reaching benchmark by defining a total of 40 diverse goal images across 4 different robot environments. In contrast to common RL benchmarks such as Atari [5] that require training over 50 different agents and thus enormous computational resources, our unsupervised RL benchmark only requires training 4 agents, which are then evaluated across many tasks, allowing for faster iteration time and making the research more accessible. Using this benchmark, we experimentally study the following scientific questions:

- Does the separation into explorer and achiever policies enable reaching more challenging goals than were previously possible and outperform state-of-the-art approaches?
- How does forward-looking exploration of goals compare to previous goal exploration strategies?
- How does the distance function affect the ability to reach goals in different types of environments?
- Can we train one general MaxNet to control different robots across visually distinct environments?

## 2 Explore Achieve Network (ExaNet)

Our aim is to build an agent that can achieve arbitrary user-specified goals after learning in the environment without any supervision. This problem presents two challenges, collecting trajectories that contain diverse goals and learning to reach these goals when specified as a goal image. We introduce a simple solution based on a world model and imagination training that addresses both challenges. The world model represents the agent's current knowledge about the environment and is used for training two policies, the explorer and the achiever. To explore novel situations, we construct an estimate of which states the world model is still uncertain about. To reach goals, we train the goal-conditioned achiever in imagination, using the images found so far as unsupervised goals. At test time, the agent reaches user-specified goals by deploying the achiever. A summary of the training procedure is given in Algorithm 1.

### 2.1 World Model

To efficiently predict potential outcomes of future actions in environments with high-dimensional image inputs, we leverage a Recurrent State Space Model (RSSM) [22] that learns to predict forward using compact model states that facilitate planning [6, 49]. In contrast to predicting forward in image

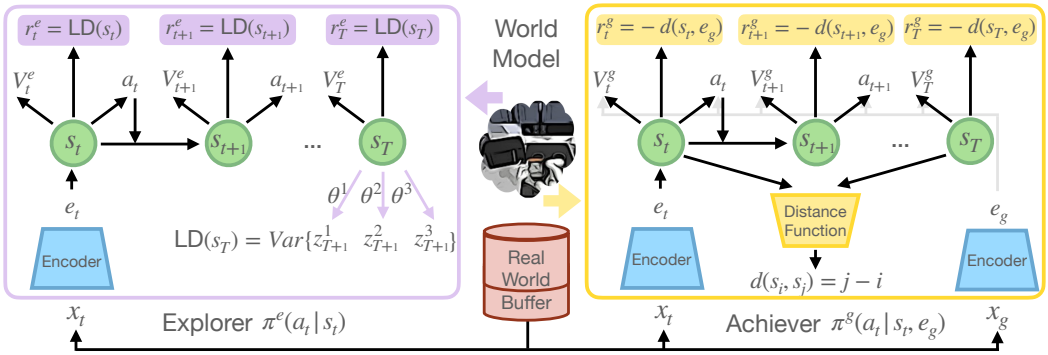


Figure 2: Explore Achieve Networks learn a single general world model that is used to train an explorer and a goal achiever policy. The explorer ($\pi^e$, left) is trained on imagined latent state rollouts of the world model $s_{t:T}$ to maximize the disagreement objective $r^e_t = \text{LD}(s_t)$. The goal achiever ($\pi^g$, right) is conditioned on a goal $g$ and is also trained on imagined rollouts to minimize a distance function $d(s_{t+1}, e_g)$. Goals are sampled randomly from replay buffer images. For training a temporal distance function, we use the imagined rollouts of the achiever and predict the number of time steps between each two states. By combining forward-looking exploration and data-efficient training of the goal achiever in imagination, Explore Achieve Network provides a simple and powerful solution for unsupervised reinforcement learning.

space, the model states enables efficient parallel planning with a large batch size and can reduce accumulating errors [38]. The world model consists of the following components:

$$\begin{aligned}
&\text{Posterior:} && q_\phi(s_t \mid s_{t-1}, a_{t-1}, e_t), \quad e_t = \text{enc}_\phi(x_t) \\
&\text{Dynamics:} && p_\phi(s_t \mid s_{t-1}, a_{t-1}) \\
&\text{Image decoder:} && p_\phi(x_t \mid s_t) \\
&\text{Reward predictor:} && p_\phi(r_t \mid s_t)
\end{aligned} \tag{1}$$

The model states $s_t$ contain a deterministic component $h_t$ and a stochastic component $z_t$ with diagonal Gaussian distribution. The deterministic component is implemented as the recurrent state of a Gated Recurrent Unit (GRU) [9]. The encoder and decoder are convolutional neural networks (CNNs) and the remaining components are multi-layer perceptrons (MLPs). The world model is trained end-to-end by optimizing the evidence lower bound (ELBO) via stochastic backpropagation [29, 37] with the Adam optimizer [28].

### 2.2 Explorer

To efficiently explore, we seek out surprising states imagined by the world model [40–42, 44], as opposed to retrospectively exploring by revisiting previously novel states [4, 7, 33]. As the world model can predict model states that correspond to unseen situations in the environment, the imagined trajectories contain more novel goals, compared to model-free exploration that is limited to the replay buffer. To collect informative novel trajectories in the environment, we train an exploration policy $\pi^e$ from the model states $s_t$ in imagination of the world model to maximize an exploration reward:

$$\text{Explorer:} \quad \pi^e(a_t \mid s_t) \qquad \text{Value:} \quad v^e(s_t) \tag{2}$$

To explore the most informative model states, we estimate the epistemic uncertainty as a disagreement of an ensemble of transition functions. We train an ensemble of 1-step models to predict the next model state from the current model state. The ensemble model is trained alongside the world model on model states produced by the encoder $q_\phi$. Because the ensemble models are initialized at random, they will differ, especially for inputs that they have not been trained on [30, 35]:

$$\text{Ensemble:} \quad f(s_t, \theta^k) = \hat{z}^k_{t+1} \quad \text{for} \quad k = 1..K \tag{3}$$

Leveraging the ensemble, we estimate the epistemic uncertainty as the ensemble disagreement. The exploration reward is the variance of the ensemble predictions and average across dimension of the model state, which approximates the expected information gain [3, 41]:

$$r^{\text{e}}_t(s_t) \doteq \frac{1}{N} \sum_n \text{Var}_{\{\text{k}\}} \left[ f(s_t, \theta_k) \right]_n \tag{4}$$

**Algorithm 1:** Explore Achieve Network (ExaNet)

**initialize:** World model $\mathcal{M}$, Replay buffer $\mathcal{D}$, Explorer $\pi^{\mathrm{e}}(a_t \mid z_t)$, Achiever $\pi^{\mathrm{g}}(a_t \mid z_t, g)$
**while** exploring **do**
  Train $\mathcal{M}$ on $\mathcal{D}$
  Train $\pi^{\mathrm{e}}$ in imagination of $\mathcal{M}$ to maximize exploration rewards $\sum_t r_t^{\mathrm{e}}$.
  Train $\pi^{\mathrm{g}}$ in imagination of $\mathcal{M}$ to maximize $\sum_t r_t^{\mathrm{g}}(z_t, g)$ for images $g \sim \mathcal{D}$.
  (Optional) Train $d(z_i, z_j)$ to predict distances $j - i$ on the imagination data from last step.
  Deploy $\pi^{\mathrm{e}}$ in the environment to explore and grow $\mathcal{D}$.
  Deploy $\pi^{\mathrm{g}}$ in the environment to achieve a goal image $g \sim \mathcal{D}$ to grow $\mathcal{D}$.
**end while**
**while** evaluating **do**
  **given:** Evaluation goal $g$
  Deploy $\pi^{\mathrm{g}}$ in the world to reach $g$.
**end while**

The explorer $\pi^e$ maximizes the sum of future exploration rewards $r_t^e$ using the Dreamer algorithm [23], which considers long-term rewards into the future by maximizing $\lambda$-returns under a learned value function. As a result, the explorer is trained to seek out situations are as informative as possible from imagined latent trajectories of the world model, and is periodically deployed in the environment to add novel trajectories to the replay buffer, so the world model and goal achiever policy can improve.

### 2.3 Achiever

To leverage the knowledge obtained by exploration for learning to reach goals, we train a goal achiever policy $\pi^g$ that receives a model state and a goal as input. Our aim is to train a general policy that is capable of reaching many diverse goals. To achieve this in a data-efficient way, it is crucial that environment trajectories that were collected with one goal in mind are reused to also learn how to reach other goals. While prior work addressed this by goal relabeling which makes off-policy policy optimization a necessity [2], we instead reuse and amplify past trajectories via the world model that is trained on past trajectories lets us generate an unlimited amount of new imagined trajectories for training the goal achiever on-policy in imagination. This simplifies policy optimization and can improve stability, while still sharing all collected experience across many goals.

$$\text{Achiever:} \qquad \pi^g(a_t \mid s_t, e_g) \qquad \text{Value:} \qquad v^g(s_t, e_g) \tag{5}$$

To train the goal achiever, we sample a goal image $x_g$ from the replay buffer and compute its embedding $e_g = \mathrm{enc}_\phi(x_g)$. The achiever aims to maximize an unsupervised goal-reaching reward $r^g(s_t, e_g)$. We discuss different choices for this reward in Section 2.4. We again use the Dreamer algorithm [23] for training, where now the value function also receives the goal embedding as input.

In addition to imagination training, we found it important to perform practice trials with the goal achiever in the true environment, so that any model inaccuracies along the goal reaching trajectories may be corrected. To perform practice trials, we sample a goal from the replay buffer and execute the goal achiever policy for that goal in the environment. These trials are interleaved with exploration episodes collected by the exploration policy in equal proportion. We note that the goal achiever learning is entirely unsupervised because the practice goals are simply images the agent encountered through exploration or during previous practice trails.

### 2.4 Latent Distances

Training the achiever policy requires us to define a goal achievement reward $r^g(s_t, e_g)$ that measures how close the latent state $s_t$ should be considered to the goal $e_g$. One simple measure is the cosine distance in the latent space obtained by inputting image observations into the world-model. However, such a distance function brings "visually" similar states together even if they could be farther apart in "temporal" manner as measured by actions needed to reach from one to other. This bias makes this suitable only to scenarios where most of pixels in the observations are directly controllable, e.g., trying to arrange robot's body in certain shape, such as RoboYoga poses in Figure 3. However, many environments contain agent as well as the world, such as manipulation involves interacting with objects that are not directly controllable. The cosine distance would try matching the entire goal image, and thus places a large weight on both matching the robot and object positions with the desired

goal. Since the robot position is directly controllable it is much easier to match, but this metric overly focuses on it, yielding poor policies that ignore objects. We address this is by using the number of timesteps it takes to move from one image to another as a distance measure [25, 26]. This ignores large changes in robot position, since these can be completed in very few steps, and will instead focus more on the objects. This temporal cost function can be learned purely in imagination rollouts from our world model allowing as much data as needed without taking any steps in the real world.

**Cosine Distance** To use cosine distance with ExaNets, for a latent state $s_t$, and a goal embedding $e^g$, we use the latent inference network $q$ to infer $s^g$, and define the reward as the cosine similarity:

$$r_t^g(s_t, e_g) \doteq \sum_i \overline{s}_{ti}\overline{s}_{gi}, \quad \text{where} \quad \overline{s}_t = s_t/\|s_t\|_2, \quad \overline{s}_g = s_g/\|s_g\|_2 \tag{6}$$

This metric is the cosine of the angle between the two vectors $s_t, s_g$ in the $N-$dimensional latent space. Using this simple metric in our Explore-Achieve framework, we obtain an effective agent for unsupervised goal reaching, especially for environments with a single controllable agent.

**Temporal Distance** To use temporal distances with ExaNets, we train a neural network $d$ to predict the number of time steps between two image predicted embeddings of an imagination trajectory. We sample a trajectory containing $z_i$ by running the current goal-reaching policy in imagination using a random initial and goal images sampled from the replay buffer. We select the second state $z_j$ to be a random state later in the same trajectory and regress towards the ground truth number of time steps between two states. We implement the temporal distance in terms of predicted image embeddings $\hat{e}_{t+k}$ in order to remove extra recurrent information:

$$\begin{aligned} &\text{Predicted CNN embedding:} && \text{emb}(s_t) = \hat{e}_t \approx e_t \\ &\text{Temporal distance:} && d_\omega(\hat{e}_t, \hat{e}_{t+k}) \approx k/H, \end{aligned} \tag{7}$$

where $H$ is the maximum distance equal to the imagination horizon. Training distance function only on imagination data from the same trajectory would cause it to predict poor distance to far away states coming from other trajectories, such as images that are impossible to reach during one episode. In order to incorporate learning signal from such far-away goals, we include them by sampling images from a different trajectory. We annotate these negative samples with the maximum possible distance, so that the agent always prefers images that were seen in the same trajectory.

$$r_t^g(s_t, e_g) = -d_\omega(\hat{e}_t, e_g), \quad \text{where} \quad \hat{e}_t = \text{emb}(s_t), \quad e_g = \text{enc}_\phi(x_g) \tag{8}$$

We note this learned distance function depends on the training data policy. However, as the policy becomes more competent, the distance estimates will be closer to the optimal number of time steps to reach a particular goal [25]. In our method, we always use the data from the latest policy to train the distance function via the imagination training, ensuring that the convergence is fast.

# 3 Experiments

We compare ExaNet to several goal reaching approaches to put its empirical performance into context and to understand the contributions of the individual components. Because not many prior methods have shown success on reaching diverse goals from image inputs, we perform an apples-to-apples comparison by implementing the baselines on top of the same world model that our method uses:

- **DDL** Dynamic Distance Learning Hartikainen et al. [25] trains a temporal distance function similar to our method. Following the original algorithm, DDL uses greedy exploration and trains the distance function on the replay buffer instead of in imagination.
- **DIAYN** Diversity is All You Need [14] learns a latent skill space and uses mutual information between skills and reach states as the objective. We augment DIAYN with our explorer policy and train a learned skill predictor to obtain a skill for a given test image [10].
- **GCSL** Goal-Conditioned Supervised Learning [19] trains the goal policy on replay buffer goals and mimics the actions that previously led to the goal. We also augment GCSL with our explorer policy, as we found no learning success without it.
- **SkewFit** SkewFit [36] uses model-free hindsight experience replay and explores by sampling goals from the latent space of a variational autoencoder [29, 37]. Being one of the state-of-the-art agents, we use the original implementation that does not use a world model or explorer policy.

We introduce three benchmarks for unsupervised goal reaching by defining goal images for a diverse set of four existing environments, shown in Figure 3:

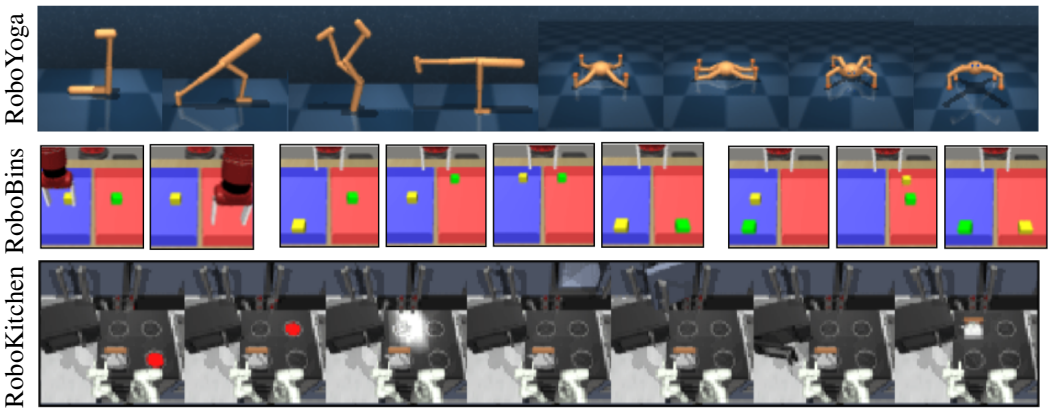


Figure 3: We benchmark our method across four visual control tasks of varying degrees of difficulty: RoboYoga (Walker, Quadruped), RoboBins, and RoboKitchen. A representative sample of the used testing goals is shown here. The RoboYoga benchmark features complex locomotion and precise control of high-dimensional agents, the RoboBins benchmark features extended manipulation tasks with manipulating multiple objects at the same time, while RoboKitchen features a variety of diverse tasks that require complex control strategies such as opening a cabinet.

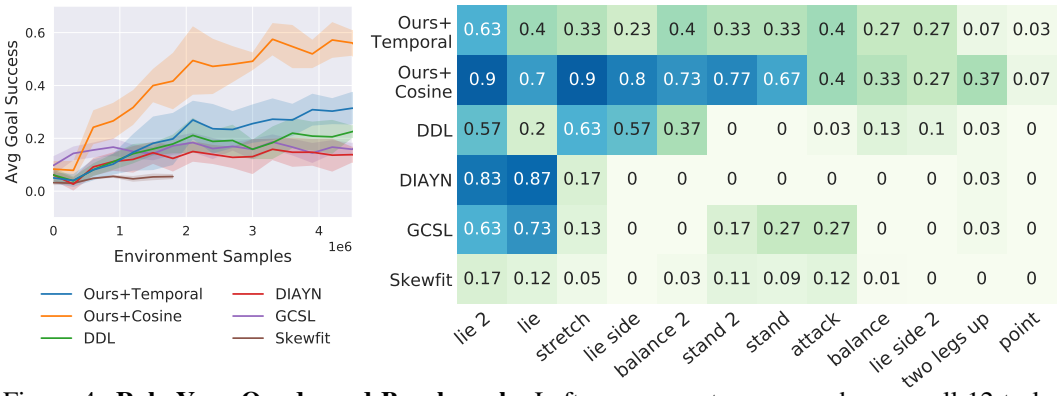


Figure 4: **RoboYoga Quadruped Benchmark.** Left: success rates averaged across all 12 tasks. Right: final performance on each specific task, ranging from light green (0) to dark blue (100%). We observe that the simple latent cosine distance function works well on this task, substantially outperforming other competing agents. In the heatmap, most agents can solve the easy tasks, but only ExaNet makes progress on solving a majority of the tasks and achieves good performance.

- **RoboYoga** We use the walker and quadruped domains of the DeepMind Control Suite [46] to defined the RoboYoga benchmark, consisting of 12 goal images that correspond to different body poses for each of the two environments, such as lying down, standing up, and balancing.
- **RoboBins** Based on MetaWorld [51], we create a scene with a Sawyer robotic arm, two bins, and two blocks of different colors. The goal images specify tasks that include reaching, manipulating only one block, and manipulating both blocks.
- **RoboKitchen** The last benchmark involves the challenging kitchen environment from [21], where a franka robot can interact with various objects including a burner, light switch, sliding cabinet, hinge cabinet, microwave, or kettle. The goal images we include describe tasks that require interacting with only one object, as well as interacting with two objects.

### 3.1 RoboYoga Benchmark

The environments in this benchmark are directly controllable, since they only contain the robot and no other objects. We recall that for such settings we expect the cosine distance to be a very effective metric, since success requires exactly matching the goal image. Training is thus faster compared to using learned temporal distances, where the metric is learned from scratch.

From Figure 4 we see that this is indeed the case for the Quadruped environment, and ExaNet with the cosine metric outperforms all prior approaches on both the Walker and Quadruped environments.

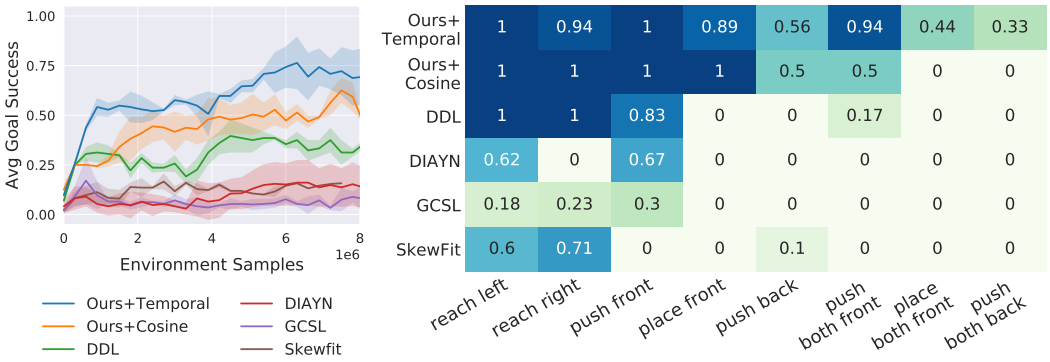


Figure 5: **RoboBin Goal Benchmark.** Left: success rates averaged across all 8 tasks. Right: final performance on each specific task, ranging from light green (0) to dark blue (100%). We observe that while the simple latent cosine distance function works on simple goals, temporal distance functions outperform it on the more challenging tasks requiring manipulations of several blocks (last three columns in the heatmap), as this distance metric is able to focus on the part of the environment that's hardest to manipulate. We further observe that other competing agents perform poorly and only solve the easiest reaching tasks, struggling either with exploration or learning the downstream policy.

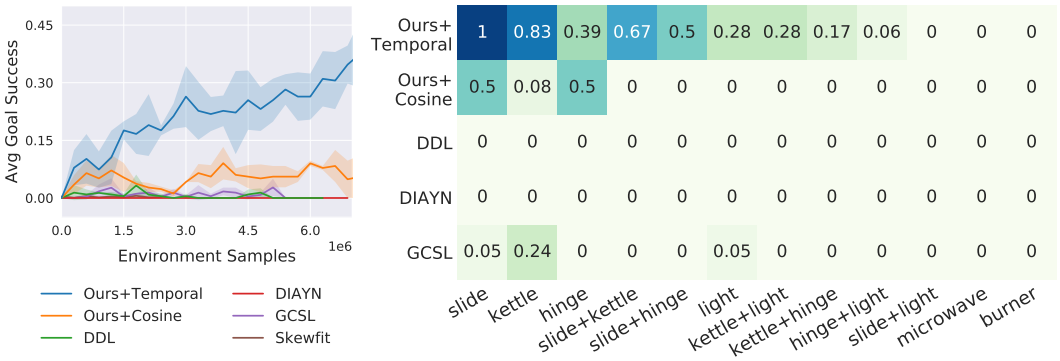


Figure 6: **RoboKitchen Benchmark.** Left: success rates averaged across all 12 tasks. Right: final performance on each specific task, ranging from light green (0) to dark blue (100%). This enviroment presents both extremely challenging exploration and downstream control, with most prior agents never solving any tasks. In contrast, ExaNet is able to learn both an effective explorer and achiever policy. Temporal distance functions help ExaNet focus on small parts such as the light switch, necessary to solve these tasks. ExaNet makes progress on four out of six base tasks, and is even able to solve combined goal images requiring e.g. both moving the kettle and opening a cabinet.

Furthermore with temporal distances ExaNet makes better progress compared to prior work on a much larger number of goals as can be seen from the per-task performance, even though average success over goals looks similar to that of DDL. We found similar results on the Walker environment, and include them in the appendix.

### 3.2 RoboBins Benchmark

This environment nvolves interaction with block objects, and so is not directly controllable, and so we expect ExaNets to perform better with the temporal distance metric. We see that ExaNets beats all prior approaches, and is also able to make progress on all goals in the benchmark.From the per-task performance in Figure 5 we see that the main difference in performance between using temporal distance and cosine can be seen in the tasks involving two blocks, which are the most complex tasks in this environment (the last 3 columns of the heatmap). The best performing prior method seems to be dynamical distances which completely solves reaching, and is about 3 times worse on pushing. This is because the the method sees a lot less data relevant to the harder tasks owing to poorer exploration since it doesn't have the disagreement objective. We see that while skewfit make some progress on reaching, it completely fails on harder tasks involving manipulation.

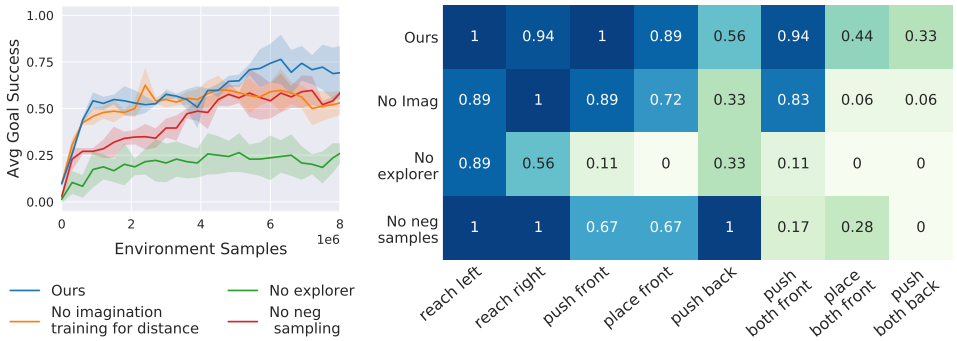


Figure 8: **Ablations** Testing different components of ExaNets with temporal distances on the RoboBins benchmark. We observe that training a separate exploration policy is crucial for solving most tasks, as the agent never discovers them in the *no explorer* version. Training temporal distance on negative samplies significantly speeds up learning, and both negative sampling and training in imagination as opposed to real data is important for performance on the hardest tasks.

### 3.3 RoboKitchen Benchmark

This benchmark involoves a very diverse set of objects, that require different manipulation behavior. We see from Figure 6 that our approach with learned temporal distance is able to learn multiple RoboKitchen tasks, some of which require sequentially completing 2 tasks in the environment. All prior methods barely make progress due to the challenging nature of this benchmark, and furthermore using the cosine distance function makes very limited progress. The gap in performance between using the two distance functions is much larger in this environment compared to RoboBins since there are many more objects and they are not as clearly visible as the blocks.

### 3.4 Single Agent Across All Environments

In the previous sections we have shown that our approach can reach diverse goals in different environments. However, we trained a new agent for every new environment, which doesn't scale well to large numbers of environments. Thus we investigate if we can train a train a single agent across all the environments in the benchmark (i.e RoboKitchen, RoboBins, Walker and Quadruped). From Figure 7 we see that our approach with learned temporal distance is able to make progress on tasks from RoboKitchen, RoboBins Reaching, RoboBins Pick & Place and Walker, while the best prior method on the single-environment tasks (DDL) mainly solves walker tasks and reaching from RoboBin.

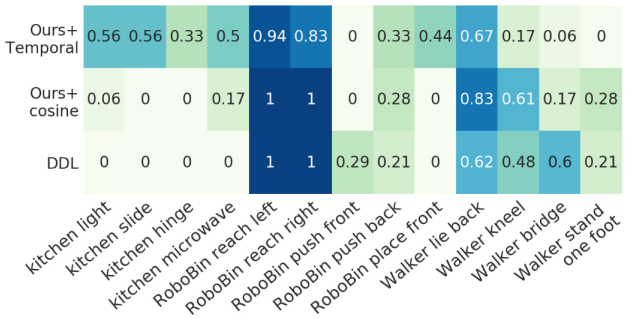


Figure 7: **Single agent** trained across Kitchen, RoboBin, Walker, with final performance on each specific task, ranging from light green (0) to dark blue (100%). ExaNet with temporal distance is able to make progress on tasks from all environments, while ExaNet+cosine and DDL don't make progress on the kitchen tasks.

### 3.5 Ablation Study

**Exploration Quality** Effective exploration is a necessary condition for goal reaching, since without diverse data it is not possible to train a good controller. In order to examine the diversity of the data collected during training, we also log the 'coincidental success', i.e whether in the process of collecting data, the agent exhibits behavior that is need to solve an evaluation task. Figure is in appendix: we see that while dynamical distances encounter reaching goals, our method encounters the harder tasks like pick and place much more often.

**Ablation of different components** We ran ablations of our approach on the RoboBins environment, where we examined the effect of removing negative sampling while training the distance function, removing disagreement for the exploration policy, and training the distance function with real world

data from the replay buffer instead of imagination data. From the plots in Figure 8 we see that disagreement for exploration is the most critical component, as without this the agent does not collect good data from which to learn from. Without negative sampling learning occurs a lot slower, and this is probably because the distance function doesn't produce reasonable outputs when queried on images that are more than horizon length apart, since it is never trained on such data. Training the distance function with real data converges to slightly lower success than using imagination data, since real data is sampled in an off-policy manner due to its limited quantity.

## 4 Related Work

The problem of learning to reach many different goals has been commonly addressed with model-free methods that learn a single goal-conditioned policy [2, 26, 39]. Recent work has combined these approaches with various ways to generate training goals, such as asymmetric self-play [32, 43] or by sampling goals of intermediate difficulty [13, 17]. These approaches can achieve remarkable performance in simulated robotic domains, however, they focus on the settings where the agent can directly perceive the low-dimensional environment state. An alternative approach to reinforcement learning altogether is to use self-imitation learning by noting that any trajectory reaches its final state and is therefore a good solution for reaching this final state [11, 19]. This approach is often a competitive baseline, however, it is sensitive to the quality of the training data as imitating very stochastic behavior may not easily produce a good policy.

A few works have attempted to scale these model-free methods to visual goals by using contrastive [48] or reconstructive [31, 36] representation learning. However, these approaches struggle to perform meaningful exploration problem as no clear reward signal is available to guide the agent toward solving interesting tasks. Chebotar et al. [8], Tian et al. [47] avoid this challenge by using a large dataset of interesting behaviors. Pong et al. [36], Zhang et al. [52] attempt to explore by generating goals similar to those that have already been seen, but do not try to explore truly novel states. A particularly relevant set of approaches used model-based methods to learn to reach goals via explicit planning [12, 16] or learning model-regularized policies [34]. However, these approaches are limited by short planning horizons. In contrast, we learn long-horizon goal-conditioned value functions which allows us to solve more challenging tasks.

A crucial challenge in visual goal-reaching settings is the choice of the reward or the cost function for the goal achieving policy. Several approaches use representation learning to create a distance in the feature space [31, 48, 50]. However, this naive distance function may not be most reflective of how hard a particular goal is to reach. One line of research has proposed using the mutual information between the current state and the goal as the distance metric [1, 10, 14, 20], however, it remains to be seen whether this approach can scale to more complex tasks.

Other approaches have attempted to learn temporal distances that measure the amount of time necessary to reach the goal. These distances can be learned with approximate dynamic programming in a Q-learning like formulation [15, 18, 26]. Our approach is most similar to that by Hartikainen et al. [25], who show that a temporal distance function can be learned via simple supervised learning on a set of recent policy experience. However, in contrast to Hartikainen et al. [25], we always train this distance function on the most recent policy data in imagination.

## 5 Conclusion

We presented Explorer Achiever Network (ExaNet), a unified agent for unsupervised reinforcement learning that is able to explore its environment and learn effective goal-reaching policies, and solves image-based tasks at test time in zero-shot way. By leveraging forward-looking exploration our agent can explore significantly better than prior approaches and discover meaningful behaviors in much more diverse environments than considered by prior work. Further, using imagination training and learned distance functions ExaNet is able to solve challenging downstream tasks, even being able to train a single policy in several different environments together. By proposing a challenging benchmark and the first agent to achieve meaningful performance on these tasks, we hope to stimulate future research on unsupervised agents, which we believe are fundamentally more scalable than traditional agents that require designing a reward or tasks during learning.

Many challenges still remain to be addressed for building effective unsupervised agents. Many of the tasks in our proposed benchmark are still largely unsolved and there remains room for progress on the algorithmic side both for the world model and for the policy optimization parts. We believe our framework provides a solid foundation for future work to make concrete progress on these goals.

# References


[1] J. Achiam, H. Edwards, D. Amodei, and P. Abbeel. Variational option discovery algorithms. *arXiv preprint arXiv:1807.10299*, 2018. 9

[2] M. Andrychowicz, F. Wolski, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, P. Abbeel, and W. Zaremba. Hindsight experience replay. *arXiv preprint arXiv:1707.01495*, 2017. 2, 4, 9

[3] P. Ball, J. Parker-Holder, A. Pacchiano, K. Choromanski, and S. Roberts. Ready policy one: World building through active learning. In *International Conference on Machine Learning*, pages 591–601. PMLR, 2020. 3

[4] M. Bellemare, S. Srinivasan, G. Ostrovski, T. Schaul, D. Saxton, and R. Munos. Unifying count-based exploration and intrinsic motivation. In *Advances in Neural Information Processing Systems*, pages 1471–1479, 2016. 3

[5] M. G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013. 2

[6] L. Buesing, T. Weber, S. Racaniere, S. Eslami, D. Rezende, D. P. Reichert, F. Viola, F. Besse, K. Gregor, D. Hassabis, et al. Learning and querying fast generative models for reinforcement learning. *arXiv preprint arXiv:1802.03006*, 2018. 2

[7] Y. Burda, H. Edwards, A. Storkey, and O. Klimov. Exploration by random network distillation. *arXiv preprint arXiv:1810.12894*, 2018. 3

[8] Y. Chebotar, K. Hausman, Y. Lu, T. Xiao, D. Kalashnikov, J. Varley, A. Irpan, B. Eysenbach, R. Julian, C. Finn, et al. Actionable models: Unsupervised offline reinforcement learning of robotic skills. *arXiv preprint arXiv:2104.07749*, 2021. 9

[9] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014. 3

[10] J. Choi, A. Sharma, S. Levine, H. Lee, and S. S. Gu. Variational empowerment as representation learning for goal-based reinforcement learning. In *Deep Reinforcement Learning workshop at the Conference on Neural Information Processing Systems (DRL)*, 2020. 5, 9

[11] Y. Ding, C. Florensa, M. Phielipp, and P. Abbeel. Goal-conditioned imitation learning. *arXiv preprint arXiv:1906.05838*, 2019. 9

[12] F. Ebert, C. Finn, S. Dasari, A. Xie, A. Lee, and S. Levine. Visual foresight: Model-based deep reinforcement learning for vision-based robotic control. *arXiv preprint arXiv:1812.00568*, 2018. 9

[13] A. Ecoffet, J. Huizinga, J. Lehman, K. O. Stanley, and J. Clune. Go-explore: a new approach for hard-exploration problems. *arXiv preprint arXiv:1901.10995*, 2019. 2, 9

[14] B. Eysenbach, A. Gupta, J. Ibarz, and S. Levine. Diversity is all you need: learning skills without a reward function. *arXiv preprint arXiv:1802.06070*, 2018. 5, 9

[15] B. Eysenbach, R. Salakhutdinov, and S. Levine. Search on the replay buffer: Bridging planning and reinforcement learning. *arXiv preprint arXiv:1906.05253*, 2019. 9

[16] C. Finn and S. Levine. Deep visual foresight for planning robot motion. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on*, pages 2786–2793. IEEE, 2017. 9

[17] C. Florensa, D. Held, X. Geng, and P. Abbeel. Automatic goal generation for reinforcement learning agents. In *International conference on machine learning*, pages 1515–1528. PMLR, 2018. 9

[18] C. Florensa, J. Degrave, N. Heess, J. T. Springenberg, and M. Riedmiller. Self-supervised learning of image embedding for continuous control. *arXiv preprint arXiv:1901.00943*, 2019. 9

[19] D. Ghosh, A. Gupta, J. Fu, A. Reddy, C. Devin, B. Eysenbach, and S. Levine. Learning to reach goals without reinforcement learning. *arXiv preprint arXiv:1912.06088*, 2019. 5, 9

[20] K. Gregor, D. J. Rezende, and D. Wierstra. Variational intrinsic control. *arXiv preprint arXiv:1611.07507*, 2016. 9

[21] A. Gupta, V. Kumar, C. Lynch, S. Levine, and K. Hausman. Relay policy learning: Solving long-horizon tasks via imitation and reinforcement learning. *arXiv preprint arXiv:1910.11956*, 2019. 6

[22] D. Hafner, T. Lillicrap, I. Fischer, R. Villegas, D. Ha, H. Lee, and J. Davidson. Learning latent dynamics for planning from pixels. *arXiv preprint arXiv:1811.04551*, 2018. 2

[23] D. Hafner, T. Lillicrap, J. Ba, and M. Norouzi. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019. 2, 4

[24] D. Hafner, T. Lillicrap, M. Norouzi, and J. Ba. Mastering atari with discrete world models. *arXiv preprint arXiv:2010.02193*, 2020. 2

[25] K. Hartikainen, X. Geng, T. Haarnoja, and S. Levine. Dynamical distance learning for semi-supervised and unsupervised skill discovery. *ICLR*, 2020. 5, 9

[26] L. P. Kaelbling. Learning to achieve goals. In *IJCAI*, pages 1094–1099. Citeseer, 1993. 5, 9

[27] L. Kaiser, M. Babaeizadeh, P. Milos, B. Osinski, R. H. Campbell, K. Czechowski, D. Erhan, C. Finn, P. Kozakowski, S. Levine, et al. Model-based reinforcement learning for atari. *arXiv preprint arXiv:1903.00374*, 2019. 2

[28] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 3

[29] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 3, 5

[30] B. Lakshminarayanan, A. Pritzel, and C. Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *arXiv preprint arXiv:1612.01474*, 2016. 3

[31] A. V. Nair, V. Pong, M. Dalal, S. Bahl, S. Lin, and S. Levine. Visual reinforcement learning with imagined goals. In *Advances in Neural Information Processing Systems*, pages 9191–9200, 2018. 2, 9

[32] O. OpenAI, M. Plappert, R. Sampedro, T. Xu, I. Akkaya, V. Kosaraju, P. Welinder, R. D'Sa, A. Petron, H. P. d. O. Pinto, et al. Asymmetric self-play for automatic goal discovery in robotic manipulation. *arXiv preprint arXiv:2101.04882*, 2021. 9

[33] D. Pathak, P. Agrawal, A. A. Efros, and T. 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. 3

[34] D. Pathak, P. Mahmoudieh, G. Luo, P. Agrawal, D. Chen, Y. Shentu, E. Shelhamer, J. Malik, A. A. Efros, and T. Darrell. Zero-shot visual imitation. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 2050–2053, 2018. 9

[35] D. Pathak, D. Gandhi, and A. Gupta. Self-supervised exploration via disagreement. In *International Conference on Machine Learning*, pages 5062–5071, 2019. 3

[36] V. H. Pong, M. Dalal, S. Lin, A. Nair, S. Bahl, and S. Levine. Skew-fit: State-covering self-supervised reinforcement learning. *arXiv preprint arXiv:1903.03698*, 2019. 2, 5, 9

[37] D. J. Rezende, S. Mohamed, and D. Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014. 3, 5

[38] V. Saxena, J. Ba, and D. Hafner. Clockwork variational autoencoders. *arXiv preprint arXiv:2102.09532*, 2021. 3

[39] T. Schaul, D. Horgan, K. Gregor, and D. Silver. Universal value function approximators. In *International conference on machine learning*, pages 1312–1320. PMLR, 2015. 9

[40] J. Schmidhuber. Curious model-building control systems. In *[Proceedings] 1991 IEEE International Joint Conference on Neural Networks*, pages 1458–1463. IEEE, 1991. 3

[41] R. Sekar, O. Rybkin, K. Daniilidis, P. Abbeel, D. Hafner, and D. Pathak. Planning to explore via self-supervised world models. *arXiv preprint arXiv:2005.05960*, 2020. 2, 3

[42] P. Shyam, W. Jaśkowski, and F. Gomez. Model-based active exploration. *arXiv preprint arXiv:1810.12162*, 2018. 2, 3

[43] S. Sukhbaatar, Z. Lin, I. Kostrikov, G. Synnaeve, A. Szlam, and R. Fergus. Intrinsic motivation and automatic curricula via asymmetric self-play. *arXiv preprint arXiv:1703.05407*, 2017. 9

[44] Y. Sun, F. Gomez, and J. Schmidhuber. Planning to be surprised: Optimal bayesian exploration in dynamic environments. In *International Conference on Artificial General Intelligence*, pages 41–51. Springer, 2011. 3

[45] R. S. Sutton. Dyna, an integrated architecture for learning, planning, and reacting. *ACM SIGART Bulletin*, 2(4):160–163, 1991. 2

[46] Y. Tassa, Y. Doron, A. Muldal, T. Erez, Y. Li, D. d. L. Casas, D. Budden, A. Abdolmaleki, J. Merel, A. Lefrancq, et al. Deepmind control suite. *arXiv preprint arXiv:1801.00690*, 2018. 6

[47] S. Tian, S. Nair, F. Ebert, S. Dasari, B. Eysenbach, C. Finn, and S. Levine. Model-based visual planning with self-supervised functional distances. *arXiv preprint arXiv:2012.15373*, 2020. 9

[48] D. Warde-Farley, T. Van de Wiele, T. Kulkarni, C. Ionescu, S. Hansen, and V. Mnih. Unsupervised control through non-parametric discriminative rewards. *arXiv preprint arXiv:1811.11359*, 2018. 2, 9

[49] M. Watter, J. Springenberg, J. Boedecker, and M. Riedmiller. Embed to control: A locally linear latent dynamics model for control from raw images. In *Advances in neural information processing systems*, pages 2746–2754, 2015. 2

[50] M. Watter, J. T. Springenberg, J. Boedecker, and M. Riedmiller. Embed to control: A locally linear latent dynamics model for control from raw images. *arXiv preprint arXiv:1506.07365*, 2015. 9

[51] T. Yu, D. Quillen, Z. He, R. Julian, K. Hausman, C. Finn, and S. Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on Robot Learning*, pages 1094–1100. PMLR, 2020. 6

[52] Y. Zhang, P. Abbeel, and L. Pinto. Automatic curriculum learning through value disagreement. *Advances in Neural Information Processing Systems*, 33, 2020. 2, 9

# 6 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]
   (c) Did you discuss any potential negative societal impacts of your work? [No] Potential societal impacts of robot learning are at this stage largely hypothetical since the technology is not yet in widespread production.
   (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]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (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)? [Yes]
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]
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (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? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]