
Hallucinative Topological Memory for Zero-Shot Visual Planning

Kara Liu^{*1} Thanard Kurutach^{*1} Christine Tung¹ Pieter Abbeel¹ Aviv Tamar^{1,2}

Abstract

In visual planning (VP), an agent learns to plan goal-directed behavior from observations of a dynamical system obtained offline, e.g., images obtained from self-supervised robot interaction. Most previous works on VP approached the problem by planning in a learned latent space, resulting in low-quality visual plans, and difficult training algorithms. Here, instead, we propose a simple VP method that plans directly in image space and displays competitive performance. We build on the semi-parametric topological memory (SPTM) method: image samples are treated as nodes in a graph, the graph connectivity is learned from image sequence data, and planning can be performed using conventional graph search methods. We propose two modifications on SPTM. First, we train an energy-based graph connectivity function using contrastive predictive coding that admits stable training. Second, to allow zero-shot planning in new domains, we learn a conditional VAE model that generates images given a context describing the domain, and use these *hallucinated* samples for building the connectivity graph and planning. We show that this simple approach significantly outperform the SOTA VP methods, in terms of both plan interpretability and success rate when using the plan to guide a trajectory-following controller. Interestingly, our method can pick up non-trivial visual properties of objects, such as their geometry, and account for it in the plans.

1. Introduction

We are interested in goal-directed planning problems where the state observations are high-dimensional images, the sys-

^{*}Equal contribution ¹Berkeley AI Research, University of California, Berkeley ²Technion. Correspondence to: Kara Liu <karamarieliu@berkeley.edu>, Thanard Kurutach <thanard.kurutach@berkeley.edu>.

tem dynamics are not known, and only a data set of state transitions is available. In particular, given a starting state observation and a goal state observation, we wish to generate a sequence of actions that transition the system from start to goal. One application for such problems is in self-supervised robot learning, where it is relatively easy to acquire such data by letting the robot explore its environment randomly, and the problem becomes how to process this data for solving various tasks (Nair et al., 2017; Wang et al., 2019; Pinto & Gupta, 2016; Finn & Levine, 2017).

Given a reward function, deep reinforcement learning (RL) can plan with high-dimensional inputs, and batch off-policy RL algorithms (Lange et al., 2012) can be applied to the problem above (Haarnoja et al., 2018; Ebert et al., 2018; Mnih et al., 2015; Schulman et al., 2015). However, goal-based planning is a sparse reward task, which is known to be difficult for RL (Andrychowicz et al., 2017). Moreover, RL provides black-box decision policies, which are not interpretable, and can only be evaluated by running them on a robot. Addressing both data-driven modeling and interpretability, the *visual planning* (VP) paradigm seeks to first generate a visual plan – a sequence of images that transition the system from start to goal, which can be understood by a human observer – and only then take actions that follow the plan using visual servoing methods.

Bearing similarity to model-based RL (Sutton et al., 1998), most VP approaches learn a low-dimensional latent variable model for the system dynamics, and plan a state-to-goal sequence by searching in the latent space (Kurutach et al., 2018; Asai, 2019; Ebert et al., 2018; Hafner et al., 2018; Nair & Finn, 2019). There are two shortcomings to this approach: training deep generative models with a structured latent space can be tricky in practice (Watter et al., 2015; Kurutach et al., 2018), and consequentially, the resulting visual plans are often of low visual fidelity.

In this work, we propose a simple VP method that plans directly in image space. We build on the semi-parametric topological memory (SPTM) method proposed by Savinov et al. (Savinov et al., 2018). In SPTM, images collected offline are treated as nodes in a graph and represent the possible states of the system. To connect nodes in this graph, an image classifier is trained to predict whether pairs of images were ‘close’ in the data or not, effectively learning

which image transitions are feasible in a small number of steps. The SPTM graph can then be used to generate a visual plan – a sequence of images between a pair of start and goal images – by directly searching the graph. SPTM has several advantages, such as producing highly interpretable visual plans and the ability to plan long-horizon behavior. Here, we ask – is such a simple scheme competitive with VP methods that plan in latent space?

To answer this question, we need to address a limitation of SPTM compared to VP methods such as visual foresight (Finn & Levine, 2017; Ebert et al., 2018). Since SPTM builds the visual plan directly from images in the data, when the environment changes – for example, the lighting varies, the camera is slightly moved, or other objects are displaced – SPTM requires *recollecting* images in the new environment; in this sense, SPTM *does not generalize in a zero-shot sense*. To tackle this issue, we assume that the environment is described using some context vector, which can be an image of the domain or any other observation data that contains enough information to extract a plan (see Figure 1 top left). We then train a conditional generative model that hallucinates possible states of the domain conditioned on the context vector. Thus, given an unseen context, the generative model hallucinates exploration data without requiring actual exploration. Using the hallucinated images, we can then perform planning in image space.

Additionally, similar to (Eysenbach et al., 2019), we find that training the graph connectivity classifier as originally proposed by (Savinov et al., 2018) requires extensive manual tuning. We replace the vanilla classifier used in SPTM with an energy-based model that employs a contrastive loss. We show that this alteration drastically improves planning robustness and quality. Finally, for planning, instead of connecting nodes in the graph according to an arbitrary threshold of the connectivity classifier, as in SPTM, we cast the planning as an inference problem, and efficiently search for the shortest path in a graph with weights proportional to the inverse of a proximity score from our energy model. Empirically, we demonstrate that this provides much smoother plans and barely requires any hyperparameter tuning. We term our approach Hallucinative Topological Memory (HTM). A visual overview of our algorithm is presented in Figure 1.

We evaluate our method on a set of simulated VP problems that require non-myopic planning, and accounting for non-trivial object properties, such as geometry, in the plans. In contrast with prior work, which only focused on the success of the method in executing a task, here we also measure the *interpretability* of visual planning, through mean opinion scores of features such as image fidelity and feasibility of the image sequence. In both measures, HTM outperforms state-of-the-art data-driven approaches such as visual foresight (Ebert et al., 2018)

and the original SPTM. The codebase and videos can be found at <https://sites.google.com/view/hallucinativetopologicalmemory>.

2. Preliminaries

Context-Conditional Visual Planning and Acting (VPA)

Problem. We consider the context-conditional visual planning problem from (Kurutach et al., 2018; Wang et al., 2019). Consider deterministic and fully-observable environments $\mathcal{E}_1, \dots, \mathcal{E}_N$ that are sampled from an environment distribution $P_{\mathcal{E}}$. Each environment \mathcal{E}_i can be described by a context vector c_i that entirely defines the dynamics $o_{t+1}^i = m(o_t^i, a_t^i | c_i)$, where o_t^i, a_t^i are the observations and actions, respectively, at timestep t from context c_i . For example, in the illustration in Figure 1, the context could represent an image of the obstacle positions, which is enough to predict the possible movement of objects in the domain.¹ As is typical in VP problems, we assume our data $\mathcal{D} = \{o_1^i, a_1^i, \dots, o_{T_i}^i, c_i\}_{i \in \{1, \dots, N\}}$ is collected in a self-supervised manner², and that in each environment \mathcal{E}_i , the observation distribution is defined as $P_o(\cdot | c_i)$.

At test time, we are presented with a new environment, its corresponding context vector c , and a pair of start and goal observations o_{start}, o_{goal} . Our goal is to use the training data to build a planner $\mathcal{K}_h(o_{start}, o_{goal}, c)$ and an h-horizon policy $\pi_h(o_{current}, o_{target})$. The planner’s task is to generate a sequence of observations between o_{start} and o_{goal} , in which any two consecutive observations are reachable within h time steps. The policy outputs an action that brings the current image to the target image within h steps which can be used to follow the generated plan. This requires both the planner and the policy to work together in zero-shot.

In this work, we first evaluate the planner and policy separately – the planner by measuring the fidelity of its plans, and the policy, by measuring its success rate in tracking a feasible plan. We then also evaluate the combined planner+policy by measuring the total success rate of the policy applied to the planned trajectories. For simplicity we will omit the subscript h for the planner and the policy.

Semi-Parametric Topological Memory (SPTM) (Savinov et al., 2018) is a visual planning method that can be used to solve a special case of VPA, where there is only a single training environment, \mathcal{E} and no context image. SPTM builds

¹We used such a context image in our experiments. We assume that in practice it is feasible to remove the robot from the domain, making this setting relevant to applications.

²Like Nair et al. (2017) and Wang et al. (2019), we assume that random actions allow us to sufficiently explore the environment for the tasks we consider. Depending on the task, this may require some engineering of the set of possible actions, e.g., a focused pick-and-place action set in Nair et al. (2017). In our domains, basic movements and rotations sufficed.

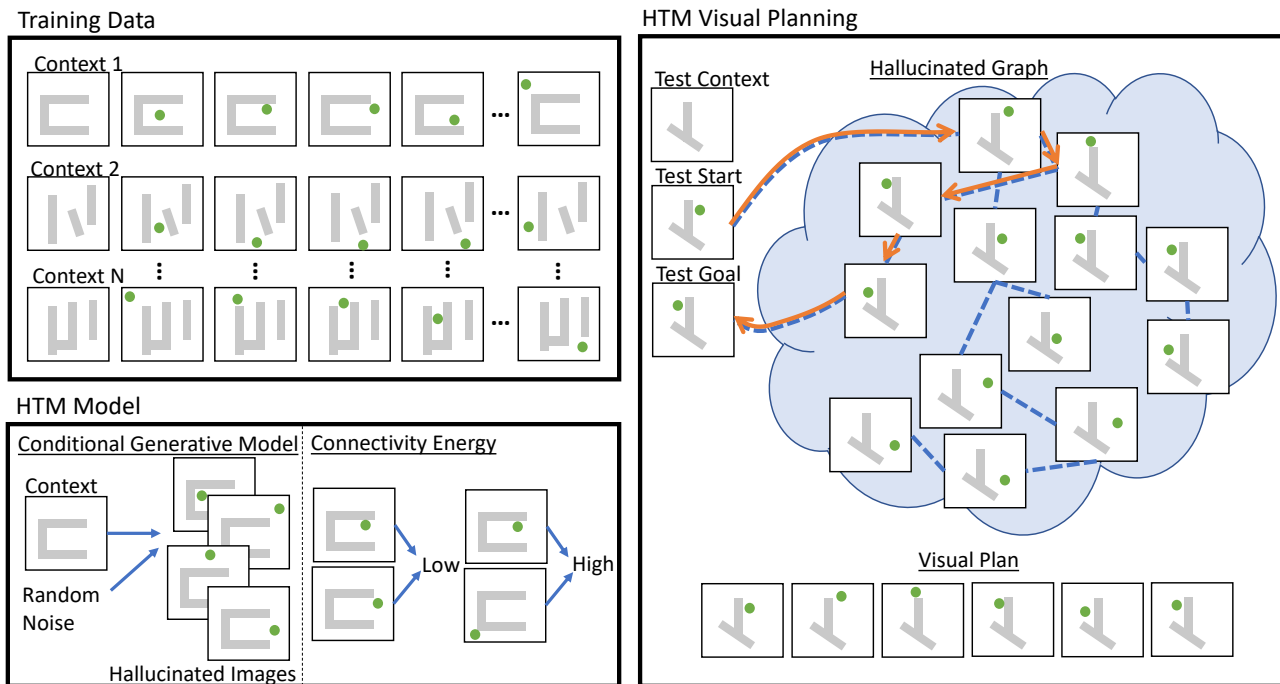


Figure 1. HTM illustration. **Top left:** data collection. In this illustration, the task is to move a green object between gray obstacles. Data consists of multiple obstacle configurations (contexts), and images of random movement of the object in each configuration. **Bottom left:** the elements of HTM. A conditional generative model is trained to hallucinate images of the object and obstacles conditioned on the obstacle image context. A connectivity energy model is trained to score pairs of images based on the feasibility of their transition. **Right:** HTM visual planning. Given a new context image and a pair of start and goal images, we first use the conditional generative model to hallucinate possible images of the object and obstacles. Then, a connectivity graph (blue dotted lines) is computed based on the connectivity energy, and we plan for the shortest path from start to goal on this graph (orange solid line). For plan execution, visual servoing is later used to track the image sequence.

a memory-based planner and an inverse-model controller. At training, a classifier R is trained to map two observation images o_i, o_j to a score $\in [0, 1]$ representing the feasibility of the transition, where images that are $\leq h$ steps apart are labeled positive and images that are $\geq l$ are negative. The policy is trained as an inverse model L , mapping a pair of observation images o_i, o_j to an appropriate action a that transitions the system from o_i to o_j .

Given an unseen environment \mathcal{E}^* , new observations are *manually* collected and organized as nodes in a graph G . Edges in the graph connect observations o_i, o_j if $R(o_i, o_j) \geq s_{shortcut}$, where $s_{shortcut}$ is a manually defined threshold. To plan, given start and goal observations o_{start} and o_{goal} , SPTM first uses R to localize, i.e., find the closest nodes in G to o_{start} and o_{goal} . A path is found by running Dijkstra’s algorithm, and the method then selects a waypoint o_{w_i} on the path which represents the farthest observation that is still feasible under R . Since both the current localized state o_i and its waypoint o_{w_i} are in the observation space, we can directly apply the inverse model and take the action a_i where $a_i = L(o_i, o_{w_i})$. After localizing to the new observation state reached by a_i , SPTM repeats the process until the

node closest to o_{goal} is reached.

Contrastive Predictive Coding (CPC) (Oord et al., 2018) is a method for learning low-dimensional representations of high-dimensional sequential data. CPC learns both an encoding of the data at every time step, and an energy function for any two observations in different time steps in suggesting their temporal correlation. A non-linear encoder g_{enc} encodes the observation o_t to a latent representation $z_t = g_{enc}(o_t)$. We maximize the mutual information between the latent representation z_t and future observation o_{t+k} with a log-bilinear model³ $f_k(o_{t+k}, o_t) = \exp(z_{t+k}^T W_k z_t)$. This model is trained to be proportional to the density ratio $p(o_{t+k}|z_t)/p(o_{t+k})$ by the CPC loss function: the cross entropy loss of correctly classifying a positive sample from a set of random observations consisting of 1 positive sample (o_t, o_{t+k}) from the paired data \mathcal{D}_{pair} and $N - 1$ negative samples (o_t, o') where o' is sampled separately from the full data \mathcal{D}_{single} :

³The original CPC model has an additional autoregressive memory variable (Oord et al., 2018). We drop it in our formulation as our domains are fully observable and do not require memory.

$$I(o_t, o_{t+1}) \geq -\mathbb{E}_{(o_t, o_{t+k}) \sim \mathcal{D}_{pair}} \left[\log \frac{f(z_t, z_{t+1})}{\sum_{\hat{o} \in \mathcal{D}_{single}} f(\hat{o}, z_t)} \right].$$

Note that the model f_k is not necessary symmetric, and therefore can capture asymmetric transition in the data. f_k can also be viewed as an inverse energy model whose outputs are high for positive samples and low for negative samples.

3. Hallucinative Topological Memory

By planning directly in image space, and composing the plan from real images (vs. planning in a learned latent space (Kurutach et al., 2018)), SPTM is guaranteed to produce high-fidelity visual plans. In addition, SPTM has been shown to solve long-horizon planning problems such as navigation from first-person view (Savinov et al., 2018). However, *SPTM is not zero-shot*: even a small change to the training environment requires collecting substantial exploration data for building the planning graph. This can be a limitation in practice, especially in robotic domains, as any interaction with the environment requires robot time, and exploring a new environment can be challenging (indeed, Savinov et al. 2018 applied manual exploration). In addition, similarly to Eysenbach et al. (2019), we found that training the connectivity classifier as proposed by Savinov et al. (2018) requires extensive hyperparameter tuning.

In this section, we propose an extension of SPTM to overcome these two challenges by employing three ideas – (1) using a conditional generative model such as CVAE (Sohn et al., 2015) or CGAN (Mirza & Osindero, 2014) to hallucinate samples in a zero-shot setting, (2) using contrastive loss for a more robust score function and planner, and (3) planning based on an approximate maximum likelihood formulation of the shortest path under uniform state distribution. We call this approach Hallucinative Topological Memory (HTM), and next detail each component in our method.

3.1. Hallucinating Samples

We propose a zero-shot learning solution for automatically building the planning graph using only a context vector of the new environment. Our idea is that, after seeing many different environments and corresponding states of the system during training, given a new environment we should be able to effectively *hallucinate* possible system states. We can then use these hallucinations in lieu of real samples from the system in order to build the planning graph. To generate images conditioned on a context, we implement a conditional generative model as depicted in Figure 1. During training, we learn the conditional distribution $p_\theta(o|c)$. During testing, when prompted with a new context vector c_{test} , we generate samples $\hat{o}_1, \dots, \hat{o}_N \sim p_\theta(o|c_{test})$ in replacement

of exploration data.

3.2. Algorithm

We now describe the HTM algorithm. Given a start observation o_{start} , a goal observation o_{goal} sampled from a potentially new environment \mathcal{E}^* , and the context vector c , we propose a 4-step planning algorithm.

1. We hallucinate exploration data $\hat{o}_1, \dots, \hat{o}_N$ by sampling from the conditional generative model $p_\theta(\cdot|c)$.
2. We build a fully-connected weighted graph $G(V, E)$ by forming connections between all generated image pairs \hat{o}_i, \hat{o}_j with learned directed edge weight w_{ij} .
3. We find the shortest path using Dijkstra’s algorithm on the learned connectivity graph G between the start and goal node.
4. We apply a local policy to follow the visual plan, attempting the next node in our shortest path for h time steps, and replan every fixed number of steps until we reach o_{goal} .

In step 2, the weights should reflect difficulty in transitioning from one state to another using a self-supervised exploration policy. The learned connectivity graph G can be viewed as a topological memory upon which we can use conventional graph planning methods to efficiently perform visual planning. In step 4, for the policy, we train an inverse model which predicts an action given the current observation and a nearby goal observation. In practice, given a transition (o_t, a_t, o_{t+1}) , we train a deep convolutional neural network $\pi(o_t, o_{t+1}) = \hat{a}_t$ to minimize the L2 loss between a_t and \hat{a}_t (Nair et al., 2017; Wang et al., 2019).

3.3. Learning the Connectivity via Contrastive Loss

A critical component in the SPTM method is the connectivity classifier that decides which image transitions are feasible. False positives may result in impossible short-cuts in the graph, while false negatives can make the plan unnecessarily long. In (Savinov et al., 2018), the classifier was trained discriminatively, using observations in the data that were reached within h steps as positive examples, and more than l steps as negative examples, where h and l are chosen arbitrarily. In practice, this leads to three important problems. First, this method is known to be sensitive to the choice of positive and negative labeling (Eysenbach et al., 2019). Second, training data are required to be long, non-cyclic trajectories for a high likelihood of sampling ‘true’ negative samples. However, self-supervised interaction data often resembles random walks that repeatedly visit a similar state, leading to inconsistent estimates on what constitutes negative data. Third, since the classifier is only trained to predict positively for temporally nearby images and negatively for temporally far away images, its predictions of *medium-distance* images can be arbitrary. This creates both

false positives and false negatives, thereby increasing short-cuts and missing edges in the graph.

To solve these problems, we propose to learn a connectivity score using contrastive predictive loss (Oord et al., 2018). We initialize a CPC encoder g_{enc} that takes in both observation and context, and a density-ratio model f_k that does not depend on the context. Through optimizing the CPC objective, f_k is trained such that positive pairs, which appear sequentially, have higher score, i.e., lower energy, than negative pairs, which are sampled randomly from the data. Thus, it serves as a proxy for the temporal distance between two observations in the sense that sequential observations should have lower energy, leading to a connectivity score for planning in the next section. Compared to the heuristic classification loss in SPTM, the CPC loss is derived from a clear objective: maximize the mutual information between the latent encodings of current and future observations. In practice, this results in less hyperparameter tuning and a smoother distance manifold in the representation space mitigating the first and the third problems.

To tackle the second problem, instead of only sampling negative data within the same trajectory as an anchor image o_t as done in SPTM, we sample any \tilde{o} that shares the same context c as o_t from the replay buffer. We also find that adding negative data sampled from $p_\theta(\cdot|c)$ can help f_k evaluate more consistently on hallucinated images. Without this trick, we find that the SPTM classifier suffers from false negatives and fails to train on short, cyclical trajectories collected by self-supervised interaction.

3.4. Edge Weight Selection

We would like edge weights to reflect the difficulty in transitioning from one state to another according causality in the data – low weight when the transition is feasible. Based on the connectivity score from the contrastive loss, we proposed two choices of computing the weight w_{ij} from node j to node i : (1) an energy model or an inverse of f_k , i.e., $1/f_k(i, j)$, and (2) a density ratio or an inverse of normalized f_k over outgoing edges from j , i.e., $\sum_{s \in V} f_k(s, j)/f_k(i, j)$. With this heuristic, the shortest path in G tries to predict reachable visual plans. In Appendix, we argue that the shortest path in graph G according to the weights in option 2 leads to maximizing trajectory likelihood bound under uniform data assumption, thus, casting planning as inference.

4. Related work

Reinforcement Learning. Most of the study of data-driven planning has been under the model-free RL framework (Schulman et al., 2015; Mnih et al., 2015; Silver et al., 2016). However, the need to design a reward function, and the fact that the learned policy does not generalize to tasks

that are not defined by the specific reward, has motivated the study of model-based approaches. Recently, Kaiser et al. (2019); Ichter & Pavone (2019) investigated model-based RL from pixels on Mujoco and Atari domains, but did not study generalization to a new environment. Finn & Levine (2017); Ebert et al. (2018) explored model-based RL with image-based goals using visual model predictive control (visual MPC). These methods rely on video prediction, and are limited in the planning horizon due to accumulating errors. In comparison, our method does not predict full trajectories but only individual images, mitigating this problem. Our method is orthogonal to and can be combined with visual MPC as a replacement for the inverse model.

Concurrently with our work, Nair & Finn (2019) propose a hierarchical visual MPC method that, similarly to our approach, optimizes a sequence of hallucinated images as sub-goals for a visual MPC controller, by searching over the latent space of a CVAE. To compute a search update over a proposed plan, the algorithm evaluates the video prediction score between consecutive subgoals making it more expensive and challenging to optimize as the number of subgoals increases. In practice, it is shown to work with the maximum of 2 subgoals. We also find that it takes approximately to 2 hours to plan a single task comparing to a few seconds in HTM making the algorithm impractical to evaluate without access to a large GPU cluster and especially in a closed-loop setting.

Self-supervised learning. Several studies investigate planning goal directed behavior from data obtained offline, e.g., by self-supervised robot interaction (Agrawal et al., 2016; Pinto & Gupta, 2016). Nair et al. (2017) use an inverse model to reach local sub-goals, but require human demonstrations of long-horizon plans. Wang et al. (Wang et al., 2019) solve the visual planning problem using a conditional version of Causal InfoGAN (Kurutach et al., 2018). Our work is not tied to specific types of generative models and in the experiments we opted for the CVAE-based approach for stability and robustness.

Classical planning and representation learning. Studies that bridge between classical planning and representation learning include (Kurutach et al., 2018; Asai & Fukunaga, 2018; Asai, 2019; Eysenbach et al., 2019). These works, however, do not consider zero-shot generalization. While Srinivas et al. (2018) and Qureshi et al. (2019) learn representations that allow goal-directed planning to unseen environments, they require expert training trajectories. Ichter & Pavone (2019) also generalizes motion planning to new environments, but require a collision checker and valid samples from test environments.

5. Experiments

We evaluate HTM on a suite of simulated tasks inspired by robotic manipulation domains. We note that recent work in visual planning (e.g., Kurutach et al. 2018; Wang et al. 2019; Ebert et al. 2018) focused on real robotic tasks with visual input. While impressive, such results can be difficult to reproduce or compare. For example, it is not clear whether manipulating a rope with the PR2 robot (Wang et al., 2019) is more or less difficult than manipulating a rigid object among many visual distractors (Ebert et al., 2018). Our suite of tasks is reproducible, contains clear evaluation metrics, and our code will be made available for evaluating other algorithms in the future.

We consider four domains in varying difficulty using Mujoco simulation (Todorov et al., 2012), as seen in Figure 2:

1. **Block wall:** A green block navigates around a static red obstacle, which can vary in position.
2. **Block wall with complex obstacle:** Similar to the above, but here the wall is a 3-link object which can vary in position, joint angles, and length, making the task significantly harder.
3. **Block insertion:** Moving a blue block, which can vary in shape, through an opening.
4. **Robot manipulation:** A simulated Sawyer robot reaching and displacing a block.



Figure 2. Evaluation suite. Block wall, block wall with complex obstacle, block insertion, and robot manipulation domains.

With the first three domains, we aim to assess how well HTM can *generalize* to new environments in a zero-shot manner, by varying the position of the obstacle, the shape of the obstacle, and the shape of the object. With the fourth domain, we aim to assess whether HTM can plan temporally-extended robotic manipulation.

We ask the following questions. First, does HTM improve **visual plan quality** over state-of-the-art VP methods (Savinov et al., 2018; Ebert et al., 2018)? Second, how does HTM **execution success rate** compare to state-of-the-art VP methods? We discuss our evaluation metrics for these attributes in Section 5.1.

We compare HTM with two state-of-the-art baselines: SPTM (Savinov et al., 2018) and Visual Foresight (Ebert et al., 2018). When evaluating zero-shot generalization, SPTM requires samples from the new environment. For a

fair comparison, we use the same samples generated by the same CVAE as HTM. Thus, in this case, we only compare between the SPTM classification scores as edge weights in the graph, and our CPC-based scores.⁴

The same low-level controller is also used to follow the plans. The Visual Foresight baseline trains a video prediction model, and then performs model predictive control (MPC), which searches for an optimal action sequence using random shooting. For the random shooting, we used 3 iterations of the cross-entropy method with 200 sample sequences. The MPC acts for 10 steps and then replans, where the planning horizon T is set to 15 as in the original implementation. Experiments with different horizons yielded worse performance. We use the state-of-the-art video predictor as proposed by Lee et al. (Lee et al., 2018) and the public code provided by the authors. For evaluating trajectories in random shooting, we studied two cost functions that are suitable for our domains: pixel MSE loss and green pixel distance. The pixel MSE loss computes the pixel distance between the predicted observations and the goal image. This provides a sparse signal when the object pixels in the plan can overlap with those of the goal. We also investigate a cost function that uses prior knowledge about the task – the position of the moving green block, which is approximated by calculating the center of mass of the green pixels. As opposed to pixel MSE, the green pixel distance provides a smooth cost function which estimates the normalized distance between the estimated block positions of the predicted observations and the goal image. Note that this assumes additional domain knowledge which HTM does not require.

5.1. Evaluation Metrics

We design a set of tests that measure both qualitative and quantitative performance of an algorithm. While the quantitative tests evaluate how successful the algorithm is in solving a task, our qualitative tests provide a measure of *plan interpretability*, which is often desired in practice.

Qualitative evaluation: Visual plans can be inspected by a human to assess their quality. Since human assessment is subjective, we devised a set of questionnaires, and for each domain, we asked 5 participants to visually score 5 randomly generated plans from each model by answering the following questions: (1) *Fidelity*: Does the pixel quality of the images resemble the training data?; (2) *Feasibility*: Is each transition in the generated plan executable by a single action step?; and (3) *Completeness*: Is the goal reachable from the last image in the plan using a single action? Answers were in the range $[0,1]$, where 0 denotes *No* to the

⁴In practice, we found that exponentiating the SPTM classifier score instead of thresholding worked slightly better, without requiring tuning a threshold. We therefore report results using this method.




Figure 3. HTM plan examples (top 3 rows) and Visual Foresight plan examples (bottom 3 rows). Note Visual Foresight is unable to conduct a long-horizon plan, and thus greedily moves in the direction of the goal state using green pixel distance cost.

proposed question and 1 means *Yes*. The mean opinion scores are reported for each model.

Quantitative In addition to generating visually sensible trajectories, a planning algorithm must also be able to successfully navigate towards a predefined goal. Thus, for each domain, we selected 20 start and goal images, each with an obstacle configuration unseen during training. Success was measured by the ability to get within some L2 distance to the goal in n steps or less, where the distance threshold and n varied on the domain but was held constant across all models. A controller specified by the algorithm executed actions given an imagined trajectory, and replanning occurred every r steps. Specific details can be found in the Appendix.

5.2. Results on Block Domains

As shown in Table 1, HTM outperforms all baselines in both qualitative and quantitative measurements across the first two domains. In the simpler block wall domain, Visual Foresight only succeed with the extra domain knowledge of using the green pixel distance. In the complex obstacle domain, Visual Foresight mostly fails to find feasible plans. SPTM, on the other hand, performed poorly on both tasks, showing the importance of our CPC-based edge weights in the graph. **Perhaps the most interesting conclusion from this experiment, however, is that even such visually simple domains, which are simulated, have a single moving object, and do not contain visual distractors or lighting/textural variations, can completely baffle state-of-the-art VP algorithms.** For the complex obstacle domain, we attribute this to the non-trivial geometric information about the obstacle shape that needs to be extracted from the context and accounted for during planning. In comparison, the real-image domains of Ebert et al. (2018),




Figure 4. HTM plan and execution. The top row demonstrates a generated visual plan on an unseen block configuration, and the bottom displays the execution to follow the plan.




Figure 5. HTM plans on real data. The agent plans to grab the green block and/or go around the obstacle with goal directed planning (no reward signal) using dataset \mathcal{D} on unseen starts and goals.

which contained many distractors, did not require much information about the shape of the objects for planning a successful pushing action.

In regards to perceptual evaluation, Visual Foresight generates realistic transitions, as seen by the high participant scores for feasibility. However, the algorithm is limited in creating a visual plan within the optimal $T = 15$ timesteps consistent with that of (Ebert et al., 2018). Thus, when confronted with a challenging task of navigating around a concave shape where the number of timesteps required exceeds T , Visual Foresight fails to construct a reliable plan (see Figure 3), and thus lacks plan completeness. Conversely, SPTM is able to imagine some trajectory that will reach the goal state. However, as mentioned above and was confirmed in the perceptual scores, SPTM fails to select feasible transitions, such as imagining a trajectory where the block will jump across the wall or split into two blocks. Our approach, on the other hand, received the highest scores of fidelity, feasibility, and completeness. Finally, we show in Figure 6 the results of our two proposed improvements to SPTM in isolation. The results clearly show that a classifier using contrastive loss outperforms that which uses Binary Cross Entropy (BCE) loss, and furthermore that the inverse of the score function for edge weighting is more successful than the best tuned version of binary edge weights through thresholding – 0 means no edge connection and 1 means an edge exists.

Algorithms	Domain	Fidelity	Feasibility	Completeness	Execution Success
HTM	1	0.75 ± .09	0.88 ± .14	1.00 ± .00	95%
	2	0.96 ± .03	0.96 ± .08	0.96 ± .08	100%
SPTM with CVAE	1	0.40 ± .11	0.00 ± .00	1.00 ± .00	55%
	2	0.92 ± .07	0.00 ± .00	1.00 ± .00	30%
Visual Foresight (Ebert et al., 2018) (pixel MSE loss)	1	0.74 ± .08	0.84 ± .16	0.04 ± .08	25%
	2	0.59 ± .16	0.64 ± .21	0.00 ± .00	0%
Visual Foresight (Ebert et al., 2018) (green pixel distance)	1	0.80 ± .07	0.84 ± .16	0.04 ± .08	90%
	2	0.69 ± .14	0.56 ± .21	0.00 ± .00	35%
Inverse Model	1	-	-	-	20%
	2	-	-	-	25%

Table 1. Qualitative and quantitative evaluation for the the block wall (1) and block wall with complex obstacle (2) domains. Qualitative data also displays the 95% confidence interval.

Algorithms	Difficulty	Execution Success
HTM	Easy	100 %
	Hard	70 %
Visual Foresight	Easy	60 %
	Hard	10 %
Inverse Model	Easy	90 %
	Hard	30 %

Table 2. Quantitative evaluation for block insertion domain. Visual Foresight (Ebert et al., 2018) was trained using pixel MSE loss.

5.3. Results for Insertion and Manipulation Domains

Unlike previous block domains, we demonstrate the zero-shot generalization ability of our approach by varying the shape and volume of the moving object itself. The challenge in planning is accounting for the orientation of a novel shape when encountering obstacles, and figuring out the best angle at which to approach a narrow passageway. We emphasize that **such geometrical reasoning must be learned from the data, and must generalize to unseen shapes.**

In practice, it might be very difficult to extract a context vector describing the environment every time the environment changes. In this domain, instead, we show that conditioning on a random image in that environmental configuration is sufficient to generate possible image samples in the same configuration, and therefore lead to successful plans. At test time, we can simply use the current image as context.

For testing, we differentiated between ‘easy’ tasks (ie. block stays on the same side of the wall) and ‘hard’ tasks (ie. block must pass through the opening). Each task had 10 random start/goal locations, and all configurations were unseen. As seen in Table 2, our method is successfully able to tackle these challenges, and an example plan and execution can be found in Figure 4. While successful on the majority of the ‘easy’ tasks, Visual Foresight proved unable to plan the rotations necessary to move the block through the opening,

and thus failed on most of the ‘hard’ tasks.

In addition, we applied HTM to robotic simulation of a Sawyer robot arm as seen in Figure 5 in which the robot needs to move a the green block to the desired location around a wall. We collect 45,000 samples of random interaction when the arm holds the green block, and 5,000 samples when the arm moves without the block. Here, we do not have different contexts, but we evaluate on unseen starts and goals. We apply HTM planning ability directly on real images from the replay buffer achieving feasible plans 12 out of 14 test tasks. We find that our visual plans avoid myopic behavior by planning to going around the thin wall, and preferring to grab the block before moving to the goal.




Figure 6. Ablation study on weight functions. We show gain in using our proposed score function and weighting function comparing to those proposed in the original SPTM by examining final *average distance* to the goal state for 10 test start/goal pairs on block with complex obstacle domain (the lower the distance, the better). For the score function, we denote our proposed energy model structured with contrastive loss as *CPC* and the classifier as proposed in (Savinov et al., 2018) with BCE loss as *SPTM*. For the edge weighting function, we test the binary thresholding from the original SPTM paper, our proposed inverse of the score function, and our proposed inverse of the normalized score function.

6. Discussion

We proposed a simple visual planning method that plans directly in image space, and generalizes in a zero-shot by hallucinating possible images conditioned on a domain context. On a suite of challenging visual planning domains, we find that our method outperforms state-of-the-art methods, and is able to pick up non-trivial geometrical information about objects in the image that is crucial for planning.

Our results further suggest that combining classical planning methods with data-driven perception can be helpful for long-horizon visual planning problems, and takes another step in bridging the gap between learning and planning. In future work, we plan to combine HTM with Visual MPC for handling more complex objects, and use object-oriented planning for handling multiple objects.

7. Acknowledgement

We thank Sudeep Dasari and Frederik Ebert for sharing code for Visual Foresight baselines and helpful discussions. Tharnard Kurutach is funded by NSF Expedition grant 1730628. Kara Liu was funded by Berkeley Deep Drive. Aviv Tamar is partly funded by the Israel Science Foundation (ISF-759/19) and the Open Philanthropy Project Fund.

References

- Agrawal, P., Nair, A. V., Abbeel, P., Malik, J., and Levine, S. Learning to poke by poking: Experiential learning of intuitive physics. In *Advances in Neural Information Processing Systems*, pp. 5074–5082, 2016.
- Andrychowicz, M., Wolski, F., Ray, A., Schneider, J., Fong, R., Welinder, P., McGrew, B., Tobin, J., Abbeel, O. P., and Zaremba, W. Hindsight experience replay. In *Advances in neural information processing systems*, pp. 5048–5058, 2017.
- Asai, M. Unsupervised grounding of plannable first-order logic representation from images. *arXiv preprint arXiv:1902.08093*, 2019.
- Asai, M. and Fukunaga, A. Classical planning in deep latent space: Bridging the subsymbolic-symbolic boundary. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Ebert, F., Finn, C., Dasari, S., Xie, A., Lee, A., and Levine, S. Visual foresight: Model-based deep reinforcement learning for vision-based robotic control. *arXiv preprint arXiv:1812.00568*, 2018.
- Eysenbach, B., Salakhutdinov, R., and Levine, S. Search on the replay buffer: Bridging planning and reinforcement learning. *arXiv preprint arXiv:1906.05253*, 2019.
- Finn, C. and Levine, S. Deep visual foresight for planning robot motion. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on*, pp. 2786–2793. IEEE, 2017.
- Haarnoja, T., Zhou, A., Hartikainen, K., Tucker, G., Ha, S., Tan, J., Kumar, V., Zhu, H., Gupta, A., Abbeel, P., et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.
- Hafner, D., Lillicrap, T., Fischer, I., Villegas, R., Ha, D., Lee, H., and Davidson, J. Learning latent dynamics for planning from pixels. *arXiv preprint arXiv:1811.04551*, 2018.
- Ichler, B. and Pavone, M. Robot motion planning in learned latent spaces. *IEEE Robotics and Automation Letters*, 4(3):2407–2414, 2019.
- Kaiser, L., Babaeizadeh, M., Milos, P., Osinski, B., Campbell, R. H., Czechowski, K., Erhan, D., Finn, C., Koza-kowski, P., Levine, S., et al. Model-based reinforcement learning for atari. *arXiv preprint arXiv:1903.00374*, 2019.
- Kurutach, T., Tamar, A., Yang, G., Russell, S. J., and Abbeel, P. Learning plannable representations with causal infogan. In *Advances in Neural Information Processing Systems*, pp. 8733–8744, 2018.

- Lange, S., Gabel, T., and Riedmiller, M. Batch reinforcement learning. In *Reinforcement learning*, pp. 45–73. Springer, 2012.
- Lee, A. X., Zhang, R., Ebert, F., Abbeel, P., Finn, C., and Levine, S. Stochastic adversarial video prediction. *arXiv preprint arXiv:1804.01523*, 2018.
- Mirza, M. and Osindero, S. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529, 2015.
- Nair, A., Chen, D., Agrawal, P., Isola, P., Abbeel, P., Malik, J., and Levine, S. Combining self-supervised learning and imitation for vision-based rope manipulation. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on*, pp. 2146–2153. IEEE, 2017.
- Nair, S. and Finn, C. Hierarchical foresight: Self-supervised learning of long-horizon tasks via visual subgoal generation. *arXiv preprint arXiv:1909.05829*, 2019.
- Oord, A. v. d., Li, Y., and Vinyals, O. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- Pinto, L. and Gupta, A. Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours. In *2016 IEEE international conference on robotics and automation (ICRA)*, pp. 3406–3413. IEEE, 2016.
- Qureshi, A. H., Simeonov, A., Bency, M. J., and Yip, M. C. Motion planning networks. In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 2118–2124. IEEE, 2019.
- Savinov, N., Dosovitskiy, A., and Koltun, V. Semi-parametric topological memory for navigation. *arXiv preprint arXiv:1803.00653*, 2018.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. Trust region policy optimization. In *International Conference on Machine Learning*, pp. 1889–1897, 2015.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484, 2016.
- Sohn, K., Lee, H., and Yan, X. Learning structured output representation using deep conditional generative models. In *Advances in neural information processing systems*, pp. 3483–3491, 2015.
- Srinivas, A., Jabri, A., Abbeel, P., Levine, S., and Finn, C. Universal planning networks. *arXiv preprint arXiv:1804.00645*, 2018.
- Sutton, R. S., Barto, A. G., et al. *Introduction to reinforcement learning*, volume 135. MIT press Cambridge, 1998.
- Todorov, E., Erez, T., and Tassa, Y. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033. IEEE, 2012.
- Wang, A., Kurutach, T., Liu, K., Abbeel, P., and Tamar, A. Learning robotic manipulation through visual planning and acting. *arXiv preprint arXiv:1905.04411*, 2019.
- Watter, M., Springenberg, J., Boedecker, J., and Riedmiller, M. Embed to control: A locally linear latent dynamics model for control from raw images. In *Advances in neural information processing systems*, pp. 2746–2754, 2015.

A. Discriminative Models: Classifier vs. Energy model

In this section, we assume the dataset as described in VPA, $\mathcal{D} = \{o_1^i, \dots, o_{T_i}^i\}_{i=1}^n$. There are two ways of learning a model to distinguish the positive from the negative transitions.

Classifier: As noted above, SPTM first trains a classifier which distinguishes between an image pair that is within h steps apart, and the images that are far apart using random sampling. The classifier is used to localize the current image and find possible next images for planning. In essence, the classifier contains the encoder g_θ that embeds the observation x and the score function f that takes the embedding of each image and output the logit for a sigmoid function. The binary cross entropy loss of the classifier $L_{SPTM}(\theta, \psi; \mathcal{D})$ is

$$\begin{aligned} &= - \sum_{(z_t, z_{t+k}) \sim \mathcal{D}} \left(\log \frac{f_\psi(z_t, z_{t+k})}{1 + f_\psi(z_t, z_{t+k})} \right. \\ &\quad \left. + \log \frac{1}{1 + f_\psi(z_t, z_t^-)} \right) \\ &= - \sum_{(z_t, z_{t+1}) \sim \mathcal{D}} \log \left[\frac{f_\psi(z_t, z_{t+k})}{f_\psi(z_t, z_{t+k}) + \alpha_\psi^t} \right] \end{aligned}$$

where $\alpha_\psi^t = 1 + f_\psi(z_t, z_t^-) + f_\psi(z_t, z_{t+k})f_\psi(z_t, z_t^-)$, and z_t^- is a random sample from \mathcal{D} .

Energy model: Another form of discriminating the the positive transition out of negative transitions is through an energy model. Oord et al. (Oord et al., 2018) learn the embeddings of the current states that are predictive of the future states. Let g be an encoder of the input x and $z = g_\theta(x)$ be the embedding. The loss function can be described as a cross entropy loss of predicting the correct sample from $N+1$ samples which contain 1 positive sample and N negative samples $L_{CPC}(\theta, \psi; \mathcal{D})$ is

$$= - \sum_{(z_t, z_{t+k}) \sim \mathcal{D}} \log \left[\frac{f_\psi(z_t, z_{t+k})}{f_\psi(z_t, z_{t+k}) + \sum_{i=1}^N f_\psi(z_t, z_t^{i-})} \right]$$

where $f_\psi(u, v) = \exp(u^T \psi v)$ and $z_t^{1-}, \dots, z_t^{N-}$ are the random samples from \mathcal{D} .

Note that when the number of negative samples is 1 the loss function resembles the SPTM.

B. Mutual Information (MI)

This quantity measures how much knowing one variable reduces the uncertainty of the other variable. More precisely, the mutual information between two random variables X and Y can be described as

$$\begin{aligned} I(X, Y) &= H(X) - H(X|Y) = H(Y) - H(Y|X) \\ &= \mathbb{E}_{X, Y} \left[\frac{p_{X, Y}}{p_X p_Y} \right]. \end{aligned}$$

C. Planning as Inference

After training the CPC objective to convergence, we have $f_k(o_{t+k}, o_t) \propto p(o_{t+k}|o_t)/p(o_{t+k})$ (Oord et al., 2018). To estimate $p(o_{t+k}|o_t)/p(o_{t+k})$, we compute the normalizing factor $\sum_{o' \in V} f_k(o', o_t)$ for each o_t by averaging over all nodes in the graph. Therefore, our non-negative weight from o_t to o_{t+k} is defined as $\omega(o_t, o_{t+k}) = \sum_{o' \in V} f_k(o', o_t) / f_k(o_{t+k}, o_t) \approx p(o_{t+k}) / p(o_{t+k}|o_t)$.

A shortest-path planning algorithm finds T, o_0, \dots, o_T that minimizes $\sum_{t=0}^{T-1} \omega(o_t, o_{t+1})$ such that $o_0 = o_{start}, o_T = o_{goal}$. By Jensen's inequality and the Markovian property of o_0, \dots, o_T we have that, $\log \frac{1}{T} \sum_{t=0}^{T-1} \omega(o_t, o_{t+1}) \geq \frac{1}{T} \sum_{t=0}^{T-1} \log \omega(o_t, o_{t+1}) = \frac{1}{T} \sum_{t=0}^{T-1} (\log p(o_{t+1}) - \log p(o_{t+1}|o_t)) = \frac{1}{T} \sum_{t=1}^{T-1} p(o_t) - \log p(o_1, \dots, o_{T-1}|o_0 = o_{start}, o_T = o_{goal})$. Thus, since $p(o_t)$ is fixed by uniform assumption, the shortest path algorithm with proposed weight ω maximizes a lower bound on the trajectory likelihood given the start and goal states. In practice, this leads to a more stable planning approach and yields more feasible plans.

D. Block Insertion Domain

In this domain, we kept the obstacle constant and varied the agent itself. In particular, we uniformly chose from 4 to 10 units, with 6 as the holdout, and then randomly placed those units such that they resembled a contiguous shape. When applying an action, we applied a vertical and horizontal force to the middle block, and also a rotation force on the first and last unit laid down, leading to a total action space of four. As our context vector, we randomly chose any image from all trajectories with that same context, as seen in Figure 7. During testing time, we randomly generated shapes from 3, 6, and 11 units. The L2 threshold distance for success was thus the total L2 distance for all units divided by the number of units.

E. Additional Results and Hyperparameters




Figure 7. Example of observations (top) and contexts (bottom) of block insertion domain.

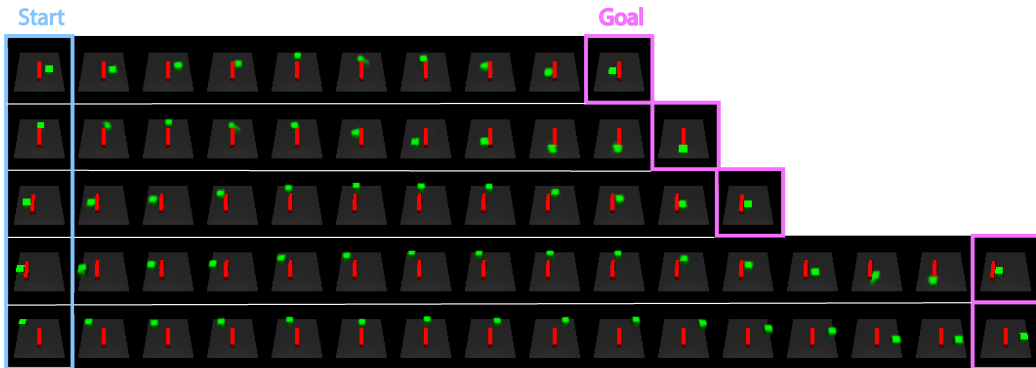


Figure 8. HTM plan examples on the block wall domain. The hallucination allows the planner to imagine how to go around the wall even though it has not seen the context before.




Figure 9. Visual Foresight plan examples on the block wall domain. The plans do not completely show the trajectory to the goal.

Table 3. Data parameters.

	Domain 1	Domain 2	Domain 3	Domain 4
no. contexts	150	400	360	1
initializations per context	50	30	20	1000
trajectory length	20	100	50	50
action space	$[-.05, .05]^2$	$[-.1, .1]^2$	$[-.05, .05]^4$	$[-1, 1]^2$
table size	2.8x2.8	2.8x2.8	.8x.8	.9x.7

Table 4. Planning hyperparameters.

	Domain 1	Domain 2	Domain 3
no. of samples from CVAE	300	500	300
L2 threshold for success (for each unit)	.5	.75	.1
n (timesteps to get to goal)	500	400	400
r (timesteps until replanning)	200	80	80