

Visual Robot Task Planning

Chris Paxton¹, Yotam Barnoy¹, Kapil Katyal¹, Raman Arora¹, Gregory D. Hager¹

Abstract—Prospection, the act of predicting the consequences of many possible futures, is intrinsic to human planning and action, and may even be at the root of consciousness. Surprisingly, this idea has been explored comparatively little in robotics. In this work, we propose a neural network architecture and associated planning algorithm that (1) learns a representation of the world useful for generating prospective futures after the application of high-level actions, (2) uses this generative model to simulate the result of sequences of high-level actions in a variety of environments, and (3) uses this same representation to evaluate these actions and perform tree search to find a sequence of high-level actions in a new environment. Models are trained via imitation learning on a variety of domains, including navigation, pick-and-place, and a surgical robotics task. Our approach allows us to visualize intermediate motion goals and learn to plan complex activity from visual information.

I. INTRODUCTION

Humans are masters at solving problems they have never encountered before. When attempting to solve a difficult problem, we are able to build a good abstract models and to picture what effects our actions will have. Some say this act — the act of prospection — is the essence of true intelligence [1]. If we want robots that can plan and act in general purpose situations just as humans do, this ability would appear crucial.

As an example, consider the task of stacking a series of colored blocks in a particular pattern, as explored in prior work [2]. A traditional planner would view this as a sequence of high-level actions, such as `pickup(block)`, `place(block, on_block)`, and so on. The planner will then decide which object gets picked up and in which order. Such tasks are often described using a formal language such as the Planning Domain Description Language (PDDL) [3]. To execute such a task on a robot, specific goal conditions and cost functions must be defined, and the preconditions and effects of each action must be specified. This is a time consuming manual undertaking [4]. Humans, on the other hand, do not require that all of this information to be given to them beforehand. We can learn models of task structure purely from observation or demonstration. We work directly with high dimensional data gathered by our senses, such as images and haptic feedback, and can reason over complex paths without being given an explicit model or structure.

Ideally, we would learn representations that could be used for all aspects of the planning problem, that also happen to be human-interpretable. For example, deep generative models such as conditional generative adversarial networks (cGANs)




Fig. 1: Example of our algorithm using learned policies to predict a good sequence of actions. Left: initial observation x_0 and current observation x_i , plus corresponding encodings h_0 and h_i . Right: predicted results of three sequential high level actions.




Fig. 2: Predicting the next step during a suturing task based on labeled surgical data. Predictions clearly show the next position of the arms.

allow us to generate realistic, interpretable future scenes [5]. In addition, a recent line of work in robotics focuses on making structured predictions to inform planning [6, 7]: So far, however, these approaches focus on making relatively short-term predictions, and do not take into account high-level variation in how a task can be performed.

We propose a supervised model that learns high-level task structure from imperfect demonstrations. Our approach then generates interpretable task plans by predicting and evaluating a sequence of high-level actions, as shown in Fig. 1. Our approach learns a pair of functions f_{enc} and f_{dec} that map into and out of a lower-dimensional hidden space \mathcal{H} , and then learns a set of other functions that operate on values in this space. Models are learned from labeled training data containing both successes and failures, are not reliant on a large number of good expert examples, and work in a number of domains including navigation, pick-and-place, and robotic suturing (Fig. 2). We also describe a planning algorithm that uses these results to simulate a set of possible futures and choose the best sequence of high-level actions to execute, resulting in realistic explorations of possible futures.

To summarize, our contributions are:

- A network architecture and training methodology for

¹Department of Computer Science, The Johns Hopkins University, Baltimore, MD, USA {cpaxton, ybarnoy1, kkatyal2, rarora8, ghager1}@jhu.edu

learning a deep representation of a planning task.

- An algorithm to employ this planning task to generate and evaluate sequences of high-level actions.
- Experiments demonstrating the model architecture and algorithm on multiple datasets.

Datasets and source code will be made available upon publication.

II. RELATED WORK

Motion Planning: In robotics, effective TAMP approaches have been developed for solving complex problems involving spatial reasoning [8]. A subset of planners focused on Partially Observed Markov Decision Process extend this capability into uncertain worlds; examples include DeSPOT, which allows manipulation of objects in cluttered and challenging scenes [9]. These methods rely on a large amount of built-in knowledge about the world, however, including object dynamics and grasp locations.

A growing number of works have explored the integration of planning and deep neural networks. For example, QMDP-nets embed learning into a planner using a combination of a filter network and a value function approximator network [10]. Similarly, value iteration networks embed a differentiable version of a planning algorithm (value iteration) into a neural network, which can then learn navigation tasks [11]. Vezhnevets proposed to generate plans as sequences of actions [12]. Other prior work employed Monte Carlo Tree Search (MCTS) together with a set of learned action and control policies for task and motion planning [13], but did not incorporate predictions.

Prediction is intrinsic to planning in a complex world. While most robotic motion planners assume a simple causal model of the world, recent work has examined learning predictive models. Lotter et al. [14] propose PredNet as a way of predicting sequences of images from sensor data, with the goal of predicting the future, and Finn et al. [7] use unsupervised learning of predictive visual models to push objects around in a plane. However, to the best of our knowledge, ours is the first work to use prospection for task planning.

Learning Generative Models.: GANs are widely considered the state of the art in learned image generation [15, 16], though they are far from the only option. The Wasserstein GAN is of particular note as an improvement over other GAN architectures [16]. Isola et al. proposed the PatchGAN, which uses an average adversarial loss over “patches” of the image, together with an L1 loss on the images as a way of training conditional GANs to produce one image from another [5].

Prior work has examined several ways of generating multiple realistic predictions [17, 15, 18]. The authors in [19] demonstrated the advantage of applying adversarial methods to imitation learning. More recently, [18] proposed to learn a deep predictive network that uses a stochastic policy over goals for manipulation tasks, but without the goal of additionally predicting the future world state.

Learning Representations for Planning: Sung et al. [20] learn a deep multimodal embedding for a variety of tasks. This representation allows them to adapt to appliances with different interfaces while reasoning over trajectories, natural language descriptions, and point clouds for a given task. Finn et al. [6] learn a deep autoencoder as a set of convolutional blocks followed by a spatial softmax; they found that this representation was useful for reinforcement learning.

III. APPROACH

We define a planning problem with continuous states $x \in \mathcal{X}$, where x contains observed information about the world such as a camera image. We augment this state with high-level actions $a \in A$ that describe the task structure. We also assume that the hidden world state $h \in \mathcal{H}$ should encode both task information (such as goals) and the underlying ground truth input from the various sensors. For example, in the block-stacking task in Fig. 1, h encodes the positions of the four blocks, the obstacle and the configuration of the arm.

Our objective is to learn a set of models representing the necessary components of this planning problem, but acting in this latent space \mathcal{H} . In other words, given a particular action a and an observed state x , we want to be able to predict both an end state x' and the optimal sequence of actions $a \in A^*$ necessary to take us there. We specifically propose that there are three components of this prediction function:

1. $f_{enc}(x) \rightarrow h \in \mathcal{H}$, a learned encoder function maps observations and descriptions to the hidden state.
2. $f_{dec}(h) \rightarrow (x)$, a decoder function that maps from the hidden state of the world to the observation space.
3. $T(h, a) \rightarrow h' \in \mathcal{H}$, the i -th learned world state transformation function, which maps to different positions in the space of possible hidden world states.

Specifically, we will first learn an action subgoal prediction function, which is a mapping $f_{dec}(T(f_{enc}(x), a)) \rightarrow (x')$. In practice, we include the hidden state of the first world observation as well in our transform function, in order to capture any information about the world that may be occluded. This gives the transform function the form $T(h_0, h, a) \rightarrow h' \in \mathcal{H}$. We assume that the hidden state h contains all the necessary information about the world to make high level decisions as long as this h_0 is available to capture change over time.

A. Model Architecture

Fig. 3 shows the architecture for visual task planning. Inputs are two images x_0 and x_i : the initial frame when the planning problem was first posed, and the current frame. We include x_0 and h_0 to capture occluded objects and changes over time from the beginning of the task. In Fig. 3, hidden states h_0, h_i, h_j, h_k are represented by averaging across channels.

Encoder and Decoder. The first training step find the transformations into and out of the learned hidden space \mathcal{H} . Specifically, we train f_{enc} and f_{dec} using the encoder-decoder architecture shown in Fig. 4. Convolutional blocks are indicated with Ck , where k is the number of filters.

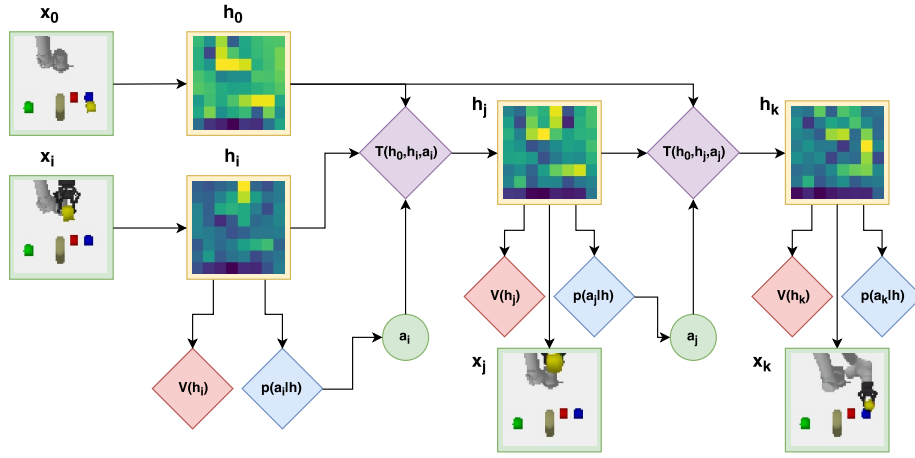


Fig. 3: Overview of the prediction network for visual task planning. We learn $f_{enc}(x)$, $f_{dec}(x)$, and $T(h, a)$ to be able to predict and visualize results of high-level actions.




Fig. 4: Encoder-decoder architecture used for learning a transform into and out of the hidden space h .

Most of our layers are 5×5 convolutions, although we used a 7×7 convolution on the first layer and we use 1×1 convolutions to project into and out of the hidden space. Each convolution is followed by an instance normalization and a ReLU activation. Stride 2 convolutions and transpose convolutions are then used to increase or decrease image size after each block. The final projection into the hidden state has a sigmoid activation in place of ReLU and is not paired with a normalization layer. In most of our examples, this hidden space is scaled down to an $8 \times 8 \times 8$ space. After every block, we add a 10% dropout layer D .

Transform function. $T(h_0, h, a)$ computes the most likely next hidden state. This function was designed to combine information about the action and two observed states, and to compute global information over the entire hidden space. We use the spatial soft argmax from previous work [21, 6, 18] Fig. 5 shows the complete transform block as used in the block stacking case study described below. For the suturing case study, with a larger input and hidden space, we add an extra set of size 64 convolutions to each side of the architecture and a corresponding skip connection, but it is otherwise the same.

Value functions. $V(h)$ computes the value of a particular hidden state h as the probability the task will be successful from that point onwards, and $Q(h_0, h, a, a')$ predicts the probability that taking action a' from the tree search node (h_0, h, a) will be successful. These are trained based on $\{0, 1\}$ labels indicating observed task success and observed failures. We also train the function $f(h_0, h, a)$ which predicts whether or not an action a completed.

Structure prior. Value functions do not necessarily indicate what happens if there are no feasible actions from a

particular state. To handle this, we learn the permissibility function $p(a'|h_0, h, a)$, which states that it is possible for a' to follow a , but does not state whether or not a' will succeed.

These last four models are trained on supervised data, but without the instance normalization in f_{enc} , f_{dec} , and T , as we saw this hurt performance. Q , p , and f were trained with two 1×1 convolutions on h and h_0 , then a concatenated, followed by $C64 - C64 - FC256 - FC128$, where FCk is a fully connected layer with k neurons. The value function $V(h)$ was a convolutional neural net of the form $C32 - C64 - C128 - FC128$.

B. Learning

We train our predictor directly on supervised pairs containing the state $x' = f_{dec}(T(f_{enc}(x), a))$ resulting from action a . First, we considered a simple L1 loss on the output images. However, this might not capture all details of complex scenes, so we also train with an augmented loss which encourages correct classification of the resulting image. This approach is in some ways similar to that used by [15], in which the authors predict images while minimizing distance in a feature space trained on a classification problem. Here, we use a combination of an L1 term and a term maximizing the cross-entropy loss on classification of the given image, which we refer to as the $L1 + \lambda C$ loss in the following, where λ is some weight. Finally, we explored using two different GAN losses: the Wasserstein GAN [16] and the pix2pix GAN from Isola et al. [5]. First we train the goal classifier $C(x)$ on labeled training data for use in testing and in training our augmented loss, then f_{enc} and f_{dec} , and finally the remaining models.

Transform Training. To encourage the model to make

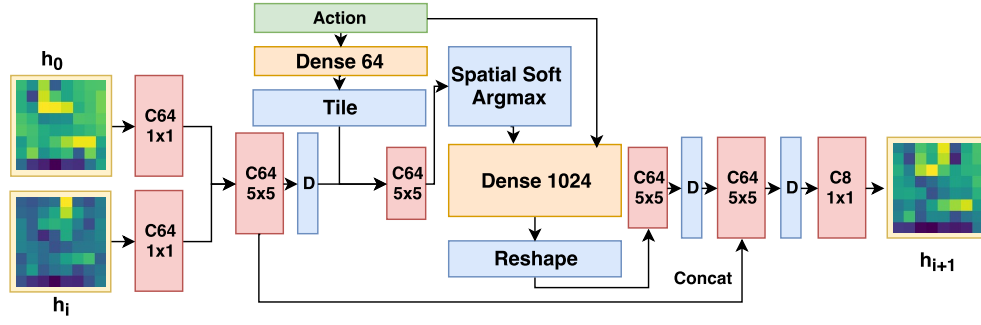


Fig. 5: Architecture of the transform function $T(h_0, h, a)$ for computing transformations to an action subgoal in the learned hidden space.

Algorithm 1 Algorithm for visual task planning with a learned state representation.

Given: max depth d , initial state x_0 , current state x , number of samples $N_{samples}$
 $h = f_{enc}(x)$, $h_0 = h$
for $i \in N_{samples}$ **do**
 EXPLORE($h_0, h, \emptyset, 0, d$)
end for
function EXPLORE(h_0, h, a, i, d)
 $v_i = \text{EVALUATE}(h_0, h, a)$
 if $i \geq d$ or $v_i < v_{failed}$ **then return** v_i **end if**
 $a' = \text{SAMPLE}(h_0, h, a)$
 $h' = T(h_0, h, a')$
 $v' = \text{EXPLORE}(h_0, h', a', i + 1, d)$
 UPDATE(a, a', v')
 return $v_i \cdot v'$
end function

predictions that remain consistent over time, we link two consecutive transforms with shared weights, and train on the sum of the L1 loss from both images, with the optional classifier loss term applied to the second image. The full training loss given ground truth predictions \hat{x}_1, \hat{x}_2 is then:

$$\mathcal{L}(x_1, x_2) = \|\hat{x}_1 - x_1\|_1 + \|\hat{x}_2 - x_2\|_1 + \lambda C(x_2)$$

We used learning rate of $1e - 4$ and 10% dropout.

C. Visual Planning with Learned Representations

We propose a variant of MCTS as a general way of exploring the tree over possible actions [13]. The general idea is that we run a loop where we repeatedly sample a possible action a' according to the learned function $Q(h_0, h, a, a')$ and use this action to simulate the effects of that high-level action $T(h_0, h, a')$. We can then execute the sequence of learned or provided black box policies to complete the motion on the robot. Each node in the tree is a unique instance of a high-level action (\emptyset for the root). The full algorithm is described in Alg. 1.

The EVALUATE function sets $v_i = V(h)$, but also checks the validity of the chosen action and determines which actions can be sampled. We also compute $f(h_0, h, a)$ to determine if the robot would successfully complete the action with some confidence c_{done} . If not this is considered a failure ($v_i = 0$). If $v' < v_{failed}$, we will halt exploration.

The SAMPLE function greedily chooses the next action a' to pursue according to a score $v(a, a')$:

$$v(a, a') = \frac{cQ(h_0, h, a, a')}{N(a, a')} + v^*(a, a')$$

where Q is the learned action-value function, $N(a, a')$ is the number of times a' was visited from a , and $v^*(a, a')$ is the best observed result from taking a' . We set $c = 10$ to encourage exploration to actions that are expected to be good. The UPDATE function is responsible for incrementing $N(a, a')$. Sampled actions a' are rejected if we predict a' is not reachable from its parent.

IV. EXPERIMENTAL SETUP

We applied the proposed method to a simulated UR5 block-stacking task implemented in PyBullet, where the robot needed to pick up a colored block and place it on top of any other colored block. The robot succeeds if it manages to stack any two blocks on top of one another and failed immediately if either it touches this obstacle or if at the end of 30 seconds the task has not been achieved. Training was performed on a relatively small number of examples: we used 6020 trials, of which 2991 were successful and 3029 were failures. The state x is a 64×64 RGB image of the scene from a fixed external camera.

We provided a set of non-optimal expert policies and randomly sampled a set of actions. This task was fairly complicated, with a total of 36 possible actions divided between two sub-tasks. Each performance was labeled a failure if either (a) it took more than 30 seconds, (b) there was a collision with the obstacle, or (c) the robot moved out of the workspace for any reason. We include both positive and negative examples in our dataset. Training was performed with Keras [22] and Tensorflow 1.5 for 45,000 iterations on an NVidia Titan Xp GPU, with a batch size of 64. Training took roughly 200 ms per batch.

A. Surgical Robot Image Prediction

Next, we explored our ability to predict the goal of the next motion on a real-world surgical robot problem. Minimally invasive surgery is a highly skilled task that requires a great deal of training; our image prediction approach could allow novice users some insight into what an expert might do in their situation. There is a growing amount of surgical




Fig. 6: Randomly sampled hidden states projected onto the manifold of the transform function $T(h_0, h, a)$.

robot video available, and a growing body of work seeks to capitalize on this to improve video prediction [23]. We used a subset of the JIGSAWS dataset to train a variant of our Visual Task Planning models on a labeled suturing task in order to predict the results of certain motions.

The JIGSAWS dataset consists of stereo video frames, with each pair labeled as belonging to one of 15 possible gestures. We used only the left frames of the video stream, and reduced the image dimensions from 640×480 to a more reasonable 96×128 . For this application, we used a slightly larger $12 \times 16 \times 8$ hidden representation.

V. RESULTS

Our models are able to generate realistic predictions of possible futures for several different tasks, and can use these predictions to make intelligent decisions about how they should move to solve planning problems. See Fig. 1 above for an example: we give the model input images x_0 and x_i , and see realistic results as it performs three actions: lifting closing the gripper, picking up the green block, and placing it on top of the blue block. In the real data set, this action failed because the robot attempted to place the green block on the red block (next to an obstacle), but here it makes a different choice and succeeds. We visualize the $8 \times 8 \times 8$ hidden layer by averaging cross the channels.

To visualize the effects of the transform function $T(h_0, h, a)$ on the learned hidden state, we randomly sampled a number of hidden states and repeatedly applied $T(h_0, h, a)$ with random actions. After 200 steps, we see results similar to those in Fig. 6, with objects and the arm positioned randomly in the scene. From left to right, these show (1) the randomly sampled hidden state, (2) a decoding of this hidden state, (3) a hidden state after 200 random transform operations, and (4) the decoded version of this hidden state.

We then performed a set of experiments to verify our model architecture, particularly in comparing different versions of the transform block T . We compared three options: the same T with the skip connection removed, and the same block with the spatial softmax and dense block replaced by a stride 2 and a stride 1 convolution to make a more traditional U-net similar to that used in prior work [5].

To compare model architectures and training strategies, we propose a simple metric: given a single frame, can we determine which action just occurred? This is computed

Model	x_1 label	x_1 error	x_2 label	x_2 error
Naive	89.3%	0.0182	77.3%	0.0276
L1	90.4%	0.0181	86.4%	0.0209
L1+0.01C	90.6%	0.0198	95.6%	0.0239
L1+0.001C	90.9%	0.0181	88.6%	0.0208
No Skips	90.1%	0.0243	89.3%	0.0271
cGAN [5]	86.5%	0.0216	79.9%	0.0260

TABLE I: Comparison of test losses as assessed by image prediction error (MAE) and image confusion.

given the same pretrained discriminator discussed in Sec. III-B. We compare versions of the loss function with and without the classifier loss term, and with this term given one of two possible weights. Both the classifier loss term and the conditional GAN discriminator term were applied to the second of two transforms, to encourage the model to generate predictions that remained consistent over time.

Table I show the results of this comparison. There were 37453 example frames from successful examples and 54077 total examples in the data set. In general, the pretrained encoder-decoder structure allowed us to reproduce high-quality images in all of our tasks. The “Naive” model indicates L1 loss with only one prediction; it performs notably worse than other models due to errors accumulating over subsequent applications of $T(h_0, h, a)$. Adding the classifier loss terms (L1+0.01C and L1+0.001C) slightly improved recognition performance looking forward in time. The cGAN was able to capture feasible texture, but often missed or made mistakes on spatial structure, and the Wasserstein GAN performed extremely poorly.

A. Plan Evaluation

Our approach is able to generate feasible action plans in unseen environments and to visualize them; see Fig. 7 for an example. The first two plans are recognized as failures, and then the algorithm correctly finds that it can pick up the red block and place it on the blue without any issues. The value function $V(h)$ correctly identified frames as coming from successful or failed trials 83.9% of the time after applying two transforms – good, considering that it is impossible to differentiate between success and failure from many frames. It correctly classified possible next actions 96.0% of the time.

We tested our method on 10 new test environments in the stacking task. On each of these environments, we performed a search with 10 samples. Our approach found 8 solutions to planning tasks executing the demonstrated high-level actions, and in 2 tasks it predicted that all of its actions would result in failures, due to proximity to the obstacle. This highlights an advantage of the visual task planning approach: in the event of a failure, the robot provides a clear explanation for why (see the second sample in Fig. 7 for an example).

B. Surgical Image Prediction

We trained our network on 36 examples in the JIGSAWS dataset, leaving out 3 for validation. We were able to generate predictions that clearly showed the location of the arms after the next gesture, as shown in Fig. 2. Image prediction created recognizable gestures, such as pulling the thread after a suture. While our results are visually impressive, error




Fig. 7: First three results from a call of the planning algorithm on a random environment.

was higher than in the stacking task: we saw mean absolute error of 0.039 and 0.062 for generated images x_1 and x_2 , respectively. This is likely because the surgical images contain a lot more subtle but functionally irrelevant data that is not fully reconstructed by our transform. It therefore looks “good enough” for human perception, but does not compare as well at a pixel-by-pixel level. In addition, there is high variability on the performance of each action and a relatively small amount of available data. Again, the cGAN did not have a measurable impact: MAE of 0.039 and 0.067 across the three test examples.

VI. CONCLUSIONS

We described an architecture for visual task planning, which learns an expressive representation that can be used to make meaningful predictions forward in time. This can be used as part of a planning algorithm that explores multiple prospective futures in order to select the best possible sequence of future actions to execute. In the future we will apply our method to real robotic examples and expand experiments on surgical data.

ACKNOWLEDGMENTS

This research project was conducted using computational resources at the Maryland Advanced Research Computing Center (MARCC). It was funded by NSF grant no. 1637949. The Titan Xp was donated by the NVIDIA Corporation.

REFERENCES

- [1] M. E. P. Seligman, P. Railton, R. F. Baumeister, and C. Sripada, “Navigating into the future or driven by the past,” *Perspectives on Psychological Science*, vol. 8, no. 2, pp. 119–141, 2013, pMID: 26172493. [Online]. Available: <https://doi.org/10.1177/1745691612474317>
- [2] D. Xu, S. Nair, Y. Zhu, J. Gao, A. Garg, L. Fei-Fei, and S. Savarese, “Neural task programming: Learning to generalize across hierarchical tasks,” *arXiv preprint arXiv:1710.01813*, 2017.
- [3] M. Ghallab, C. Knoblock, D. Wilkins, A. Barrett, D. Christianson, M. Friedman, C. Kwok, K. Golden, S. Penberthy, D. E. Smith, *et al.*, “PDDL—the planning domain definition language,” <http://www.citeulike.org/group/13785/article/4097279>, 1998.
- [4] M. Beetz, D. Jain, L. Mosenlechner, M. Tenorth, L. Kunze, N. Blodow, and D. Pangercic, “Cognition-enabled autonomous robot control for the realization of home chore task intelligence,” *Proceedings of the IEEE*, vol. 100, no. 8, pp. 2454–2471, 2012.
- [5] P. Isola, J. Zhu, T. Zhou, and A. A. Efros, “Image-to-image translation with conditional adversarial networks,” *CoRR*, vol. abs/1611.07004, 2016. [Online]. Available: <http://arxiv.org/abs/1611.07004>
- [6] C. Finn, X. Y. Tan, Y. Duan, T. Darrell, S. Levine, and P. Abbeel, “Deep spatial autoencoders for visuomotor learning,” in *ICRA*. IEEE, 2016, pp. 512–519.
- [7] C. Finn, I. Goodfellow, and S. Levine, “Unsupervised learning for physical interaction through video prediction,” in *Advances in Neural Information Processing Systems*, 2016, pp. 64–72.
- [8] M. Toussaint, “Logic-geometric programming: An optimization-based approach to combined task and motion planning,” in *IJCAI*, 2015.
- [9] J. K. Li, D. Hsu, and W. S. Lee, “Act to see and see to act: POMDP planning for objects search in clutter,” in *IROS*. IEEE, 2016, pp. 5701–5707.
- [10] P. Karkus, D. Hsu, and W. S. Lee, “Qmdp-net: Deep learning for planning under partial observability,” *arXiv preprint arXiv:1703.06692*, 2017.
- [11] A. Tamar, Y. Wu, G. Thomas, S. Levine, and P. Abbeel, “Value iteration networks,” in *Advances in Neural Information Processing Systems*, 2016, pp. 2154–2162.
- [12] A. Vezhnevets, V. Mnih, S. Osindero, A. Graves, O. Vinyals, J. Agapiou, *et al.*, “Strategic attentive writer for learning macro-actions,” in *Advances in Neural Information Processing Systems*, 2016, pp. 3486–3494.
- [13] C. Paxton, V. Raman, G. D. Hager, and M. Kobilarov, “Combining neural networks and tree search for task and motion planning in challenging environments,” 2017.
- [14] W. Lotter, G. Kreiman, and D. Cox, “Deep predictive coding networks for video prediction and unsupervised learning,” *arXiv preprint arXiv:1605.08104*, 2016.
- [15] Q. Chen and V. Koltun, “Photographic image synthesis with cascaded refinement networks,” *arXiv preprint arXiv:1707.09405*, 2017.
- [16] M. Arjovsky, S. Chintala, and L. Bottou, “Wasserstein GAN,” *stat*, vol. 1050, p. 9, 2017.
- [17] C. Rupprecht, I. Laina, R. DiPietro, M. Baust, F. Tombari, N. Navab, and G. D. Hager, “Learning in an uncertain world: Representing ambiguity through multiple hypotheses,” *ICCV*, 2017.
- [18] A. Ghadirzadeh, A. Maki, D. Kragic, and M. Björkman, “Deep predictive policy training using reinforcement learning,” 2017.
- [19] J. Ho and S. Ermon, “Generative adversarial imitation learning,” in *Advances in Neural Information Processing Systems*, 2016, pp. 4565–4573.
- [20] J. Sung, S. H. Jin, I. Lenz, and A. Saxena, “Robobarista: Learning to manipulate novel objects via deep multimodal embedding,” *arXiv preprint arXiv:1601.02705*, 2016.
- [21] S. Levine, C. Finn, T. Darrell, and P. Abbeel, “End-to-end training of deep visuomotor policies,” *Journal of Machine Learning Research*, vol. 17, no. 39, pp. 1–40, 2016.
- [22] F. Chollet, “Keras,” <https://github.com/fchollet/keras>, 2015.
- [23] S. S. Vedula, A. O. Malpani, L. Tao, G. Chen, Y. Gao, P. Poddar, N. Ahmidi, C. Paxton, R. Vidal, S. Khudanpur, *et al.*, “Analysis of the structure of surgical activity for a suturing and knot-tying task,” *PloS one*, vol. 11, no. 3, p. e0149174, 2016.