# **Procedure Planning in Instructional Videos**

 $\begin{array}{l} \mbox{Chien-Yi Chang}^{[0000-0001-6508-1161]}, \mbox{ De-An Huang}^{[0000-0002-6945-7768]}, \\ \mbox{ Danfei Xu}^{[0000-0002-8744-3861]}, \mbox{ Ehsan Adeli}^{[0000-0002-0579-7763]}, \\ \mbox{ Li Fei-Fei}^{[0000-0002-7481-0810]}, \mbox{ and Juan Carlos Niebles}^{[0000-0001-8225-9793]} \end{array}$ 

Stanford University

Abstract. In this paper, we study the problem of procedure planning in instructional videos, which can be seen as a step towards enabling autonomous agents to plan for complex tasks in everyday settings such as cooking. Given the current visual observation of the world and a visual goal, we ask the question "What actions need to be taken in order to achieve the goal?". The key technical challenge is to learn structured and plannable state and action spaces directly from unstructured videos. We address this challenge by proposing Dual Dynamics Networks (DDN), a framework that explicitly leverages the structured priors imposed by the conjugate relationships between states and actions in a learned plannable latent space. We evaluate our method on real-world instructional videos. Our experiments show that DDN learns plannable representations that lead to better planning performance compared to existing planning approaches and neural network policies.

**Keywords:** latent space planning; task planning; video understanding; representation for action and skill;

### 1 Introduction

What does it take for an autonomous agent to perform complex tasks in everyday settings, such as cooking in a kitchen? The most crucial ability is to know what actions should be taken in order to achieve the goal. In other words, the agent needs to make goal-conditioned decisions sequentially based on its perception of the environment. This sequential decision making process is often referred to as *planning* in the robotics literature. While recent works have shown great promise in learning to plan for simple tasks in structured environments such as pushing objects or stacking blocks on a table [7, 8, 31], it is unclear how these approaches will scale to visually complex, unstructured environments as seen in instructional videos [37, 38] (recordings of human performing everyday tasks).

In this paper, we study procedure planning in instructional videos, which can be seen as a step towards enabling autonomous agents to plan for complex tasks in everyday settings. Given the current visual observation of the world and a visual goal, we ask the question "What actions need to be taken in order to achieve the goal?". As illustrated in Figure 1, we define the *procedure planning* problem as: given a current visual observation  $o_t$  and a visual goal  $o_q$  that



**Fig. 1.** We study the problem of *procedure planning* in instructional videos. The goal is to generate a sequence of actions towards a desired goal such as making a mushroom omelette. Given a start observation  $o_t$  and a visual goal  $o_g$ , our model plans a sequence of actions  $\{a_i\}$  (red arrows) that can bring the start towards the goal. In addition, our model also predicts a sequence of intermediate states  $\{x_i\}$  (black dots).

indicates the desired final configuration, the model should plan a sequence of high-level actions  $\{a_t, \dots, a_{t+T-1}\}$  that can bring the underlying state of  $o_t$  to that of  $o_g$ . Here, T is the horizon of planning, which defines how many steps of task-level actions we allow the model to take.

The key technical challenge of the proposed procedure planning problem is how to learn structured and plannable state and action spaces directly from unstructured real videos. Since instructional videos are visually complex and the tasks are often described at high levels of abstractions, one can imagine indefinitely growing semantic state and action spaces from the visually complex scenes and high-level descriptions, which prevents the application of classical symbolic planning approaches [10, 15] as they require a given set of predicates for a well-defined state space.

We address this challenge by explicitly leveraging the conjugate relationships between states and actions to impose structured priors on the learned latent representations. Our key insight is that in addition to modeling the action as a transformation between states, we can also treat the state as the precondition for the next action because the state contains the history of previous actions. Following this intuition, we jointly train two modules: a forward dynamics model that captures the transitional probabilities between states and a conjugate dynamics model that utilizes the conjugate constraints for better optimization. We call our proposed method Dual Dynamics Networks (DDN), a neural network framework that learns plannable representations from instructional videos and performs procedure planning in the learned latent space.

We evaluate our approach on real-world instructional videos [38] and show that the learned latent representations and our DDN formulation are effective for planning under this setting. Our approach significantly outperforms existing planning baselines for procedure planning. We further show that our model is able to generalize to variations of the start and goal observations. In addition, we show that our approach can be applied to a related problem called walkthrough planning [17] and outperforms existing methods.

Our contributions are three-fold. First, we introduce the problem of procedure planning in instructional videos, which can be seen as a step towards enabling autonomous agents to plan for complex tasks in everyday settings such as cooking in a kitchen. Second, we propose Dual Dynamics Network (DDN), a framework that explicitly leverages structured priors imposed by the conjugate relationships between states and actions for latent space planning. Third, we evaluate our approach extensively and show that it significantly outperforms existing planning and video understanding methods for the proposed problem of procedure planning.

# 2 Related Work

**Task Planning.** Procedure planning is closely related to task planning widely studied in classical AI and robotics. Task planning is the problem of finding a sequence of task-level actions to reach the desired goal state from the current state. In the task planning literature, most studies rely on a pre-defined planning problem domain for the task [10, 15, 21]. Our work diverges from those because our proposed model can perform task-level, long-horizon planning in the visual and semantic space without requiring a hand-defined symbolic planning domain. **Planning from Pixels.** Recent works have shown that deep networks can learn to plan directly from pixel observations in domains such as table-top manipulation [17, 30, 31], navigation in VizDoom [25], and locomotion in joint space [5]. However, learning to plan from unstructured high-dimensional observations is still challenging [7,11], especially for long-horizon, complex tasks that we want to address in procedure planning. A closely related method is Universal Planning Networks (UPN) [31], which uses a gradient descent planner to learn representations from expert demonstrations. However, it assumes the action space to be differentiable. Alternatively, one can also learn the forward dynamics by optimizing the data log-likelihood from the actions [11]. We use a similar formulation and further propose the conjugate dynamics model to expedite the latent space learning. Without using explicit action supervision, causal InfoGAN [17] extracts state representations by learning salient features that describe the causal structure of toy data. In contrast to [17], our model operates directly on real-world videos and handle the semantics of actions with sequential learning.

Understanding Instructional Videos. There has been a growing interest in analyzing instructional videos [16, 23, 33, 37, 38] by studying a variety of challenging tasks. Some of the tasks ask the question "What is happening?", such as action recognition and temporal action segmentation [4, 13, 28], state understanding [2], video summarization/captioning [29, 32, 37], retrieval [32] etc. The others ask the question "What is going to happen?", such as early action

recognition [9, 35], action label prediction [1, 6, 9, 14, 22, 27, 29, 36], video prediction [18–20, 24, 26, 34], etc. However, due to the large uncertainty in human activities, the correct answer to this question is often not unique. In this paper, we take a different angle and instead ask the question: "What actions need to be taken in order to achieve a given visual goal?" Rather than requiring the model to predict potentially unbounded future actions conditioned only on history, we ask the model to infer future actions conditioned on both history and goal. Such goal-conditioned formulation immediately resolves the ambiguity when the history does not encode adequate information to bound future actions.

## 3 Method

We are interested in planning in real-world instructional videos. The key technical challenge is how to learn structured and plannable state and action spaces directly from unstructured real videos. We take a latent space approach by learning plannable representations of the visual observations and actions, along with the forward and conjugate dynamics models in the latent space. We will first define the procedure planning problem setup and how to address it using a latent space planning approach. We will then discuss how we learn the latent space and leverage the conjugate relationships between states and actions to avoid trivial solutions to our optimization. Finally, we will present the algorithms for procedure planning and walkthrough planning [17] in the learned plannable space.

### 3.1 Problem Formulation

As illustrated in Figure 1, given a current visual observation  $o_t$  and a visual goal  $o_g$  that indicates the desired final configuration, we aim to plan a sequence of actions  $\pi = \{a_t, \dots, a_{t+T-1}\}$  that can bring the underlying state of  $o_t$  to that of  $o_g$ . T is the horizon of planning, which defines how many steps of task-level actions we allow the model to take. We formulate the problem as latent space planning [17,31]. Concretely, we learn a model that can plan in some latent space spanned by the mapping functions f and g that encode the visual observation o and action a to a semantic state f(o) = x (e.g., from the observed frames to cooked eggs with mushrooms) and a latent action state  $g(a) = \bar{a}$  respectively.

Inspired by the classical Markovian Decision Process (MDP) [3], we assume that in this latent space there exists a forward dynamics  $\mathcal{T}(x_{t+1}|x_t, \bar{a}_t)$  that predicts the future state  $x_{t+1}$  given the current state  $x_t$  and the applicable action  $\bar{a}_t$ . Using the forward dynamics, we can perform sampling based planning [10] by applying different actions and search for the desired goal state  $x_g$ . Specifically, given  $f(\cdot)$ ,  $g(\cdot)$  and  $\mathcal{T}$ , we can find a plan  $\pi = \{a_t, \dots, a_{t+T-1}\}$  by (i) mapping from the visual space to the latent space  $x_t = f(o_t), x_g = f(o_g)$  and (ii) search in the latent space using  $\mathcal{T}(x_{t+1}|x_t, \bar{a}_t)$  to find the sequence of actions that can bring  $x_t$  to  $x_g$ . We will discuss details of this procedure in Section 3.3.



**Fig. 2.** (a) Our forward dynamics model  $\mathcal{T}$  predicts the next state based on the current state and action. (b) We learn the conjugate dynamics model  $\mathcal{P}$  jointly with  $\mathcal{T}$  to restrict the possible state mapping f and action embedding g. (c) At training time,  $\mathcal{T}$  takes as inputs the current state  $x_t$  and predicted current action  $\hat{a}_t$  to predict the next state  $\hat{x}_{t+1}$ , and the conjugate dynamics model  $\mathcal{P}$  takes as inputs the predicted next state  $\hat{x}_{t+1}$  and the current action  $\bar{a}_t$  to predict the next applicable action  $\hat{a}_{t+1}$ .

### 3.2 Learning Plannable Representations

In this section, we discuss how to learn the embedding functions  $f(\cdot)$ ,  $g(\cdot)$ , and forward dynamics  $\mathcal{T}$  from data. One possible approach is to directly optimize  $f(\cdot)$  with some surrogate loss function such as mutual information [17], without explicitly modeling  $\mathcal{T}(\cdot|x,\bar{a})$  and  $g(\cdot)$ . This approach is limited because it assumes a strong correspondence between the visual observation o and the semantic state x = f(o). In real-world videos, however, a small change in visual space  $\Delta o$  can induce a large variation in the semantic space  $\Delta x$  and vice versa. An alternative approach is to formulate a differentiable objective jointly with  $f(\cdot)$ ,  $g(\cdot)$ , and  $\mathcal{T}(\cdot|x,\bar{a})$  [31]. This approach also falls short when applied to procedure planning, because it often requires the action space to be continuous and differentiable, while we have an unstructured and discrete action space in real-world instructional videos.

**Forward Dynamics.** To address the above challenges, we propose to learn  $f(\cdot)$ ,  $g(\cdot)$ , and  $\mathcal{T}(\cdot|x,\bar{a})$ , jointly with a conjugate dynamics model  $\mathcal{P}$  which provides further constraints on the learned latent spaces. We consider training data  $\mathcal{D} = \{(o_t^i, a_t^i, o_{t+1}^i)\}_{i=1}^n$ , which corresponds to the triplets of current visual observation, the action to be taken, and the next visual observation. The triplets can be curated from existing instructional video datasets and the details will be discussed in Section 4. The learning problem is to find  $\mathcal{T}$  that minimizes:

$$\mathcal{L}_{T}(\mathcal{T}, f, g; \mathcal{D}) = -\frac{1}{n} \sum_{i=1}^{n} \log \Pr(x_{t+1}^{i} | x_{t}^{i}, \bar{a}_{t}^{i}; \mathcal{T})$$

$$= -\frac{1}{n} \sum_{i=1}^{n} \log \Pr(f(o_{t+1}^{i}) | f(o_{t}^{i}), g(a_{t}^{i}); \mathcal{T}).$$
(1)

The architectures are shown in Figure 2(a).

**Conjugate Dynamics.** There exists a large number of  $\mathcal{T}(\cdot|x,\bar{a})$  that can minimize  $\mathcal{L}_T$  and our model can easily overfit to the training data in  $\mathcal{D}$ . In this case, it is hard for the learned model to generalize to unseen visual observations. It is thus crucial to impose structured priors on the model. Our insight is to leverage the conjugate relationship between states and actions [12] to provide further constraints on  $f(\cdot)$  and  $q(\cdot)$  and improve the optimization of  $\mathcal{L}_T$ . Please note that this is also how our proposed method diverges from the classical MDP setting and previous works [17,31]. We can treat the loss in Eq. (1) as leveraging the standard relationship between states and actions to learn  $f(\cdot)$ : Given the current semantic state  $x_t$ , applying a latent action  $\bar{a}_t$  would bring it to a new state  $x_{t+1}$ .  $\mathcal{L}_T$  is encouraging  $f(\cdot)$  to fit  $x_{t+1} = f(o_{t+1})$  and to be consistent with the prediction of  $\mathcal{T}(\cdot|x,\bar{a})$ . On the other hand, the conjugate relationship between states and actions also implies the following: Given the previous action  $\bar{a}_{t-1}$ , the current state  $x_t$  constraints the possible actions  $\bar{a}_t$  because  $x_t$  needs to satisfy the precondition of  $\bar{a}_t$ . For example, if  $\bar{a}_{t-1}$  representing 'pour eggs to pan' is followed by  $\bar{a}_t$  representing 'cook it', then the state  $x_t$  in between must satisfy the precondition of  $\bar{a}_t$ , that is: to cook the eggs, the eggs should be in the pan. Based on this intuition, we further propose to learn a conjugate dynamics model  $\mathcal{P}(\bar{a}_t | x_t, \bar{a}_{t-1})$ , which is illustrated in Figure 2(b). The conjugate dynamics model  $\mathcal{P}$  takes as inputs the current state and the previous action to predict the current applicable action. Formally, the learning problem is to find  $\mathcal{P}$  that minimizes the following function:

$$\mathcal{L}_{P}(\mathcal{P}, f, g, h; \mathcal{D}) = -\frac{1}{n} \sum_{i=1}^{n} \left[ \log \Pr(\bar{a}_{t}^{i} | x_{t}^{i}, \bar{a}_{t-1}^{i}; \mathcal{P}) + \phi(h(g(a_{t}^{i})), a_{t}^{i}) \right]$$

$$= -\frac{1}{n} \sum_{i=1}^{n} \left[ \log \Pr\left(g(a_{t}^{i}) | f(o_{t}^{i}), g(a_{t-1}^{i}); \mathcal{P}\right) + \phi(h(g(a_{t}^{i})), a_{t}^{i}) \right].$$
(2)

where the mapping functions f and g have shared parameters across  $\mathcal{T}$  and  $\mathcal{P}$ . h is an inverse mapping function that decodes the action a from the latent action representation  $\bar{a}$ .  $\phi(\cdot, \cdot)$  measures the distance in the action space.

Now we have discussed all the components for Dual Dynamics Network (DDN), a framework for latent space task planning in instructional videos. In our framework, we encode the visual observations and the actions into latent space states and actions with  $f(\cdot)$  and  $g(\cdot)$ , and jointly learn from video demonstrations a forward dynamics model  $\mathcal{T}$  that captures the transitional probabilities and a conjugate dynamics model  $\mathcal{P}$  that utilizes the conjugate relationship between states and actions by minimizing the following combined loss function:

$$\mathcal{L}(\mathcal{T}, \mathcal{P}, f, g, g^{-1}; \mathcal{D}) = \alpha \cdot \mathcal{L}_T + \mathcal{L}_P.$$
(3)

 $\alpha$  is a weighting coefficient for the combined loss. The learning process is illustrated in Figure 2(c). During training, the forward dynamics model  $\mathcal{T}$  takes the current state  $x_t$  and predicted current action  $\hat{a}_t$  to predict the next state  $\hat{x}_{t+1}$ , and the conjugate dynamics model  $\mathcal{P}$  takes the predicted next state  $\hat{x}_{t+1}$  and the current action  $\bar{a}_t$  to predict the next applicable action  $\hat{a}_{t+1}$ .

### Algorithm 1 Procedure Planning

1: **Inputs:** Current and goal observations  $o_t$ ,  $o_g$ ; learned models  $f(\cdot)$ ,  $h(\cdot)$ ,  $\mathcal{T}$ ,  $\mathcal{P}$ ; max iteration  $\beta$ , threshold  $\epsilon$ , beam size  $\eta$ 

2:  $x \leftarrow f(o_t), x_q \leftarrow f(o_q)$  $3: x^* \leftarrow x, \, \bar{a} \leftarrow \emptyset$ 4:  $q \leftarrow PriorityQueue(), q \leftarrow q \cup \{(x, \bar{a})\}$ 5: while iteration  $<\beta$  and  $||x^* - x_g||_2^2 > \epsilon$  do 6:  $(x, \bar{a}) \leftarrow Pop(q)$  $\bar{a}' \sim \mathcal{P}(\cdot|x, \bar{a})$ 7: for  $\bar{a}_i \in \bar{a}'$  do 8:  $x \leftarrow \mathcal{T}(\cdot | x, \bar{a}_i)$ 9: 10:  $q \leftarrow q \cup \{(x, \bar{a}_i)\}$ if  $||x - x_g||_2^2 < ||x^* - x_g||_2^2$  then 11:  $x^* \leftarrow x$ 12: $q \leftarrow Sort(q)$ 13:14:  $q \leftarrow q[:\eta]$ 15:  $\{\bar{a}_i^*\} \leftarrow Backtracking(q, x^*)$ 16: return  $\{h(\bar{a}_i^*)\}$ 

### 3.3 Planning in Latent Space

In this section, we discuss how to use DDN for planning in instructional videos. The general paradigm is illustrated in Figure 3. At inference time, our full model rollouts by sampling the action from the conjugate dynamics  $\mathcal{P}$  and the next state from the forward dynamics  $\mathcal{T}$ .  $\mathcal{P}$  captures the applicable actions from the current state and improves planning performance. In the following, we first discuss how we can leverage the learned models in Section 3.2 to perform procedure planning in latent space. We then describe how our models can be applied to walkthrough planning [17], where the objective is to output the visual waypoints/subgoals between the current observation and the goal observation.

**Procedure Planning.** Using the learned models, we perform sampling-based forward planning [10] to plan a sequence of actions to achieve the goal. The process is shown in Algorithm 1. Given the current and the goal observations  $o_t$  and  $o_g$ , we first map them to the latent space with  $f(\cdot)$ :  $x_t = f(o_t), x_g = f(o_g)$ . In contrast to symbolic planning, we do not have a list of *applicable* actions to apply in the search process. One additional advantage of having jointly learned the conjugate dynamics model  $\mathcal{P}$  is that we can efficiently sample the actions to apply using  $\mathcal{P}(\cdot|x_t, \bar{a}_{t-1})$ . Based on  $\bar{a}_t$  and  $x_t$ , we can obtain  $x_{t+1}$  using  $\mathcal{T}(\cdot|x, \bar{a})$ . The search process continues for a max iteration of  $\beta$  and threshold  $\epsilon$  while maintaining a priority queue of size  $\eta$ .

Walkthrough Planning. Kurutach *et al.* proposed walkthrough planning in [17]. The outputs of walkthrough planning can serve as visual signals of the subgoals to guide task executions. In addition, it is also helpful for interpretation by visualizing what the model has learned. The details of the process are shown in Algorithm 2. Given the pool of visual observations  $\{o_i\}$ , we can first construct the score matrix  $R_{i,j}$  to capture the transition probability between two video clips

Algorithm 2 Walkthrough Planning

1: Inputs: All observations  $\{o_i\}|_{i=1}^T$ , learned models  $f(\cdot), g(\cdot), \mathcal{T}, \mathcal{P}$ , horizon T2:  $b \leftarrow \emptyset$ 3: for i in  $\{1...T\}$  do 4:  $x_i \leftarrow f(o_i)$ 5: for i in  $\{1...T\}$  do 6: for j in  $\{1...T\}$  do 7:  $R_{i,j} \leftarrow \sum_a \mathcal{T}(x_j | x_i, g(a)) \mathcal{P}(g(a) | x_i)$ 8:  $\{o_{b(i)}\} \leftarrow \arg \max_{b \in Perm(T)} \sum_{i=1}^T R_{b(i),b(i+1)}$ 9: return  $\{o_{b(i)}\}$ 





Fig. 3. At inference time, our full model rollouts by sampling the action from the conjugate dynamics  $\mathcal{P}$  and the next state from the forward dynamics  $\mathcal{T}$ .  $\mathcal{P}$  essentially captures the applicable actions and improves planning performance.

Fig. 4. Our model consistently outperforms other baselines as the horizon of planning T increases. When  $T \ge 6$ , the start observation and the visual goal are separated by 6 or more than 6 high-level actions. In this case the success rates of all baselines are less than 1%, suggesting there are more variations in the ways to reach the goal.

 $o_i$  and  $o_j$  using our learned model  $\mathcal{T}$  and  $\mathcal{P}$ . We can then perform walkthrough planning by finding the path of length T that starts at  $o_t$  and ends at  $o_g$ , while maximizing the total score. If the pool of video clips is all the clips in the same instructional video, then the problem is equivalent to finding a permutation function  $b : \{1, 2, ..., T\} \rightarrow \{1, 2, ..., T\}$  that maximizes the total score along the permutation path, under the constraints that b(1) = 1, b(T) = T.

### 4 Experiments

We aim to answer the following questions in our experiments: (i) Can we learn plannable representations from real-world instructional videos? How does DDN compare to existing latent space planning methods? (ii) How important are the forward and conjugate dynamics? (iii) Can we apply DDN to walkthrough planning problem [17] and how does it compare to existing methods?

We answer the first two questions with ablation studies on real-world instructional videos. We show how our model can generalize to various start and goal observations. For the last question, we apply our model to walkthrough planning [17] and show that our approach outperforms existing approaches.

Dataset. To train and test our proposed model, we curate a video dataset of size N that takes the form:  $\{(o_t^i, a_t^i, ..., a_{t+T-1}^i, o_{t+T}^i)\}_{i=1}^N$ . Each example contains T high-level actions, and each action can last from tens of seconds to 6 minutes. Our dataset is adapted from the 2750 labeled videos in CrossTask [38], averaging 4.57 minutes in duration, for a total of 212 hours. Each video depicts one of the 18 long-horizon tasks like Grill Steak, Make Pancakes, or Change a Tire. The videos have manually annotated temporal segmentation boundaries and action labels. Please note that since our method requires full supervision, we will leave how to utilize unlabeled instructional videos [23, 38] as future work. We encourage the reader refer to the supplementary material for data curation details. We divide the dataset into 70%/30% splits for training and testing. We use the precomputed features provided in CrossTask in all baselines and ablations for fair comparison. Each 3200-dimensional feature vector is a concatenation of the I3D, Resnet-152, and audio VGG features. The action space in CrossTask is given by enumerating all 105 combinations of the predicates and objects, which is shared across all 18 tasks. The distance function  $\phi$  in CrossTask is the cross-entropy.

**Implementation Details.** The state mapping f and action embedding g are 128-dimensional.  $\mathcal{T}$  is a 2-layer MLPs with 128 units that takes the outputs of f and g as inputs. The network  $\mathcal{P}$  shares the state and the action embedding with  $\mathcal{T}$  and we introduce recurrence to  $\mathcal{P}$  by replacing the MLPs with a twolayer RNN of 128 hidden size, and concatenate the goal embedding as input. The action decoder h takes the outputs of  $\mathcal{P}$  as inputs and outputs actions as defined by the dataset. We set hyperparameters  $\alpha = 0.001$ ,  $\beta = 20T$ ,  $\eta = 20$ , and  $\epsilon = 1e - 5$ . We use a 5-fold cross validation on the training split to set the hyper-parameters, and report performance on the test set. We train our model for 200 epochs with batch size of 256 on a single GTX 1080 Ti GPU. We use Adam optimizer with learning rate of 1e - 4 and schedule the learning rate to decay with a decay factor 0.5 and a patience epoch of 5.

#### 4.1 Evaluating Procedure Planning

In procedure planning, the inputs are the start video clip  $o_t$  and the goal video clip  $o_q$ , and the output is a sequence of T actions that brings  $o_t$  to  $o_q$ .

**Baselines.** We evaluate our model in learning plannable representations from real-world videos and compare with the following baselines and ablations:

- *Random Policy*. This baseline randomly selects an action from all actions. We include this baseline to show the empirical lower bound of performance.

- *Retrieval-Based (RB).* The procedure planning problem is formulated as a goal-conditioned decision making process. In contrast, one might approach this problem from a more static view which is analogous to [32]: Given the unseen start and goal visual observations, the retrieval-based baseline finds the nearest neighbor of the start and goal pair in the training set by minimizing the distance in the learned feature space, and then directly output the associated ground truth action labels. We use the same features as in our full model for fair comparison.

**Table 1.** Procedure planning results. Our model significantly outperforms baselines. With  $\sim 10\%$  improvement of accuracy, our model is able to improve success rate by 8 times compared to *Ours w/o T*. This shows the importance of reasoning over the full sequence. \* indicates re-implementations.

	T = 3			T = 4		
	Success Rate	Accuracy	mIoU	Success Rate	Accuracy	mIoU
Random	< 0.01%	0.94%	1.66%	< 0.01%	0.83%	1.66%
RB [32]*	8.05%	23.3%	32.06%	3.95%	22.22%	36.97%
WLTDO $[5]^*$	1.87%	21.64%	31.70%	0.77%	17.92~%	26.43%
UAAA $[6]^*$	2.15%	20.21%	30.87%	0.98%	19.86%	27.09%
UPN [31]*	2.89%	24.39%	31.56%	1.19%	21.59%	27.85%
Ours w/o $\mathcal{P}$	< 0.01%	2.61%	0.86%	< 0.01%	2.51%	1.14%
Ours w/o ${\cal T}$	1.55%	18.66%	28.81%	0.65%	15.97%	26.54%
Ours	12.18%	31.29%	47.48%	5.97%	<b>27.10</b> %	<b>48.46</b> %

- WLTDO [5]. Ehsani *et al.* proposed an action planning model for egocentric dog videos. This baseline is a recurrent model for the planning task. Given two non-consecutive observations, it predicts a sequence of actions taking the state from the first observation to the second observation. We modify the author's implementation to use the same features as our full model, and add softmax layer to output discrete actions.

- Uncertainty-Aware Anticipation of Activities (UAAA) [6]. This baseline has a two-step approach that uses RNN-HMM to infer the action labels in the observed frames, and then use an autoregressive model to predict the future action labels. We re-implement the model and modify it to condition on both the observed frames (start observations) and the visual goal.

- Universal Planning Networks (UPN) [31]. UPN is the closest to ours among existing works. Similar to our approach, UPN also aims to learn a plannable representation using supervision from the imitation loss function at training. However, it assumes a continuous and differentiable action space to enable gradient-based planning, which might not be applicable to the discrete action space in CrossTask. We re-implement UPN and adapt it to output discrete actions by adding a softmax layer.

- Ours  $w/o \mathcal{T}$ . We compare to the ablation of our model without learning the forward dynamics  $\mathcal{T}$ , where we only  $\mathcal{P}$  directly outputs the actions based on the previous action and the current state. We implement  $\mathcal{P}$  with an RNN and concatenate the goal and start as input. In this case, this ablation is equivalent to a goal-conditional RNN policy directly trained with expert actions. This ablation can also be seen as a re-implementation of the RNN-based model in [1].

- Ours  $w/o \mathcal{P}$ . We compare to the ablation without learning the conjugate dynamics. In this case, the joint optimization of the forward dynamics  $\mathcal{T}$  and the mapping  $f(\cdot)$  to the latent space can easily overfit to the training sequences.

Metrics. We use three metrics for comparison. The first is *success rate*. Although we do not have access to the underlying environment to evaluate the policies by

![](_page_10_Picture_1.jpeg)

Fig. 5. Examples of planned action sequences by DDN (ours) and the RB baseline. In the second example, DDN is not able to capture the subtle visual cues that the tool is already in the man's hand, so there is no need to get the tools out.

![](_page_10_Figure_3.jpeg)

**Fig. 6.** Examples of generalizations to unseen start/goal. (a) The 9 snapshot observations in making French toasts. (b) Our proposed model is robust to changes of start and goal observations among different stages in the video, and is able to output reasonable plans to achieve the goal and reason about the essential steps. For example, in the third row, the start  $o_3$  and goal  $o_9$  are separated by 6 actions. While the model has never seen such data during training, it successfully plans the 4 essential actions by merging repeated actions and further adding the pivotal actions of flipping and plating.

executing them in simulation, we consider a plan as a success if all the actions in the plan are the same as those in the ground truth. This is a reasonable approximation because we consider a fixed number of steps, and there is less variation in the ways to complete the task. The second metric we consider is the *accuracy* of the actions at each step, which does not require the whole sequence to match the ground truth as in the success rate metric, but only looks at each individual time step. We take the mean over the actions to balance the effect of repeating actions. The third metric we use is *mean Intersection over Union* (mIoU), which is the least strict of all the metrics we use. We compare the IoU by  $\frac{|\{a_t\} \cap \{a_t^*\}|}{|\{a_t\} \cup \{a_t^*\}|}$ , where  $\{a_t^*\}$  is the set of ground truth actions, and  $\{a_t\}$  is the set of predicted actions. We use IoU to capture the cases where the model understands what steps are required, but fails to discern the order of actions.

	<i>T</i> =	= 3	T = 4		
	Hamming	Pair Acc.	Hamming	Pair Acc.	
Random	1.06	46.85%	1.95	52.23%	
RB [32]*	0.88	56.23%	1.80	55.42%	
VO [36]*	1.02	49.06%	1.99	50.31%	
Causal InfoGAN [17]	0.57	71.55%	1.36	68.41%	
Ours w/o $\mathcal{T}$	0.99	50.45%	2.01	47.39%	
Ours w/o $\mathcal{P}$	0.33	83.33%	1.08	77.11%	
Ours	0.26	86.81%	0.88	81.21%	

**Table 2.** Results for walkthrough planning. Our model significantly outperforms the baseline by explicitly reasoning what actions need to be performed first, and is less distracted by the visual appearances. \* indicates re-implementations.

**Results.** The results are shown in Table 1. UPN is able to learn representations that perform reasonably well compared to the random baseline. However, as the action space in instructional videos is not continuous, the gradient-based planner is not able to work well. Both WLTDO and UAAA perform similar to Ours w/o  $\mathcal{T}$ , which can be seen an RNN goal-conditional policy directly trained with imitation objectives. We also note that Ours w/o  $\mathcal{P}$  cannot learn reasonable plannable representations without the conjugate dynamics. Our full model combines the strengths of planning and action imitation objective as conjugate constraints, which enables us to learn plannable representations from real-world videos to outperform all the baseline approaches on all metrics. In Figure 4, we further show our model's performance as the planning horizon increases. Our model consistently outperforms the RB baseline for all metrics because DDN explicitly imposes the structured priors using  $\mathcal{T}$  and  $\mathcal{P}$  to find the sequence of actions that reaches the goal.

Figure 5 shows some qualitative results. We note that it is difficult for the predicted sequence of actions to be exactly the same as the ground-truth, which explains why the success rates in Table 1 and Figure 4 are low in absolute value. The results are nevertheless semantically reasonable. Furthermore, as shown in Figure 6, it allows our model to generalize to different start observations when fixing the goal and vice versa. Figure 6(a) shows the 9 snapshot observations in making French toast. In the second row of Figure 6(b), we pick  $o_3$  and  $o_7$  as the start and goal observations and ask the model to plan for 4 actions to reach the goal. The model successfully recognizes that there are two French toasts and dip and pan fry the bread twice. In the third row, we change the goal to  $o_9$ . In this case, the start and goal are separated by 6 actions. While the model has never seen such data during training, it successfully plans the 4 essential actions by merging repeated actions and further adding the pivotal actions of flipping and plating. Similarly, in the first row, the model can also generalize to different start observations when fixing the goal.

### 4.2 Evaluating Intermediate States with Walkthrough Planning

Given the start and goal video clips, we have shown that our model is able to plan a sequence of actions that brings the start to the goal. At the same time, our model also predicts a sequence of intermediate states. In this section, we evaluate and visualize these predicted intermediate states. Specifically, we show how our proposed method can be apply to *walkthrough planning* [17], where the objective is to output the visual waypoints between the start and goal observations.

**Experimental Setup.** We evaluate the predicted intermediate state representations by using them to retrieve visual subgoals for task completion. In this way, the model only needs to predict lower-dimensional representations that can be used to retrieve the correct video clips from a pool of candidates. Specifically, we use all the video clips in the original video as the video clip pool. In this case, the task is equivalent to sorting the intermediate video clips while fixing the first and the last video clips.

Additional Baselines. We further compare to the following approaches for walkthrough planning:

- Visual Ordering (VO). As the task is reduced to sorting a pool of video clips [36], one baseline is to directly learn a model  $V(o_1, o_2)$  to see if  $o_1$  and  $o_2$  are consecutive video clips in the same video. Given  $V(o_1, o_2)$ , we can find the order of the candidate video clips by maximizing the total score given by V using a greedy-based search. We learned  $V(o_1, o_2)$  using the same setup as our  $\mathcal{T}$ , only that the actions a are not used as the input.

- Causal InfoGAN (CIGAN) [17]. CIGAN learns plannable representations by maximizing the mutual information between the representations and the visual observations. Additionally, the latent space is assumed to follow the forward dynamics of a certain class of actions. CIGAN is able to perform walkthrough planning using minimal supervision. We modify the author's implementation to use the same features as our full model for fair comparison.

**Metrics.** Let  $Y = (y_1, ..., y_T)$  be a sequence of the ground truth order, and  $b: \{1, 2, ..., T\} \rightarrow \{1, 2, ..., T\}$  be the permutation function such that the prediction is  $\hat{Y} = (y_{b(1)}, ..., y_{b(T)})$ . We use the following two metrics to evaluate the walkthrough planning outputs order:

- Hamming Distance: counts the number of  $\{i | i \neq b(i)\}$ .

- Pairwise Accuracy: calculates if the order between a pair i and j is respected by b(i) and b(j). It is given by  $\frac{2}{T(T-1)}\sum_{i < j, i \neq j}^{T} \{b(i) < b(j)\}.$ 

**Results.** The results are shown in Table 2. VO is unable to improve much over chance without modeling the actions. RB also struggles to perform well because it cannot handle the subsequence level variations that are not seen in the training data. By maximizing the mutual information with an adversarial loss, CIGAN [17] is able to learn reasonable models beyond Random without using action supervision. However, the complexity of the instructional videos requires explicit modeling of the forward dynamics conditioned on the semantic actions. Our full model learned for procedure planning successfully transfers to walkthrough planning and significantly outperforms all baselines on all metrics. It is interesting to see that Ours w/o  $\mathcal{P}$  actually outperforms Ours w/o  $\mathcal{T}$  in

![](_page_13_Picture_1.jpeg)

Fig. 7. Examples of visualized predicted intermediate states by DDN (ours) and two baselines. Green/red arrows indicate correct/incorrect transitions between two visual steps. Our model successfully predicts sequence of state representations with correct ordering, while the baselines are more or less distracted by the visual appearances.

this case. Ours w/o  $\mathcal{P}$  learns the forward dynamics  $\mathcal{T}$  to directly anticipate the visual effect of an action a, which is better suited for walkthrough planning. By just using  $\mathcal{P}$  the model is more accurate at predicting the actions, suggesting a trade-off between action and state space modeling. Our full model successfully combines the strength of the two and imposes the structured priors through conjugate constraints to both procedure planning and walkthrough planning.

In Figure 7, we visualize some examples of the predicted intermediates states, where the task is to change tires (left) and grill a steak (right). Our full model is able to predict a sequence of intermediate states with correct ordering. Specifically, the most challenging step in changing tire (left) is the second step, where the person goes to take the tools, and the video is visually different from the rest of the steps. Neither of the baseline models is able to understand that to perform the rest of the steps, the person needs to get the tools first.

# 5 Conclusion

We presented Dual Dynamics Networks (DDN), a framework for procedure planning in real-world instructional videos. DDN is able to learn plannable representations directly from unstructured videos by explicitly leveraging the structured priors imposed by the conjugate relationships between states and actions on the latent space. Our experimental results show our framework significantly outperforms a variety of baselines across different metrics. Our work can be seen as a step towards the goal of enabling autonomous agents to learn from real-world demonstrations and plan for complex tasks like humans. In future work, we intend to incorporate object-oriented models to further explore the objects and predicates relations in visually complex environments.

Acknowledgements. Toyota Research Institute ("TRI") provided funds to assist the authors with their research but this article solely reflects the opinions and conclusions of its authors and not TRI or any other Toyota entity.

15

# References

- 1. Abu Farha, Y., Richard, A., Gall, J.: When will you do what?-anticipating temporal occurrences of activities. In: CVPR (2018)
- 2. Alayrac, J.B., Laptev, I., Sivic, J., Lacoste-Julien, S.: Joint discovery of object states and manipulation actions. In: ICCV (2017)
- Bellman, R.: A markovian decision process. Journal of mathematics and mechanics pp. 679–684 (1957)
- Chang, C.Y., Huang, D.A., Sui, Y., Fei-Fei, L., Niebles, J.C.: D3tw: Discriminative differentiable dynamic time warping for weakly supervised action alignment and segmentation. In: CVPR (2019)
- 5. Ehsani, K., Bagherinezhad, H., Redmon, J., Mottaghi, R., Farhadi, A.: Who let the dogs out? modeling dog behavior from visual data. In: CVPR (2018)
- Farha, Y.A., Gall, J.: Uncertainty-aware anticipation of activities. arXiv preprint arXiv:1908.09540 (2019)
- Finn, C., Levine, S.: Deep visual foresight for planning robot motion. In: ICRA (2017)
- 8. Finn, C., Tan, X.Y., Duan, Y., Darrell, T., Levine, S., Abbeel, P.: Deep spatial autoencoders for visuomotor learning. In: ICRA (2016)
- Furnari, A., Farinella, G.M.: What would you expect? anticipating egocentric actions with rolling-unrolling lstms and modality attention. In: ICCV. pp. 6252–6261 (2019)
- 10. Ghallab, M., Nau, D., Traverso, P.: Automated Planning: theory and practice. Elsevier (2004)
- 11. Hafner, D., Lillicrap, T., Fischer, I., Villegas, R., Ha, D., Lee, H., Davidson, J.: Learning latent dynamics for planning from pixels. In: ICML (2019)
- 12. Hayes, B., Scassellati, B.: Autonomously constructing hierarchical task networks for planning and human-robot collaboration. In: ICRA (2016)
- Huang, D.A., Fei-Fei, L., Niebles, J.C.: Connectionist temporal modeling for weakly supervised action labeling. In: ECCV (2016)
- Ke, Q., Fritz, M., Schiele, B.: Time-conditioned action anticipation in one shot. In: CVPR (2019)
- Konidaris, G., Kaelbling, L.P., Lozano-Perez, T.: From skills to symbols: Learning symbolic representations for abstract high-level planning. Journal of Artificial Intelligence Research 61, 215–289 (2018)
- 16. Kuehne, H., Arslan, A., Serre, T.: The language of actions: Recovering the syntax and semantics of goal-directed human activities. In: CVPR (2014)
- 17. Kurutach, T., Tamar, A., Yang, G., Russell, S.J., Abbeel, P.: Learning plannable representations with causal infogan. In: NeurIPS (2018)
- Lan, T., Chen, T.C., Savarese, S.: A hierarchical representation for future action prediction. In: ECCV (2014)
- Lee, A.X., Zhang, R., Ebert, F., Abbeel, P., Finn, C., Levine, S.: Stochastic adversarial video prediction. arXiv preprint arXiv:1804.01523 (2018)
- Lu, C., Hirsch, M., Scholkopf, B.: Flexible spatio-temporal networks for video prediction. In: CVPR (2017)
- McDermott, D., Ghallab, M., Howe, A., Knoblock, C., Ram, A., Veloso, M., Weld, D., Wilkins, D.: Pddl-the planning domain definition language (1998)
- 22. Mehrasa, N., Jyothi, A.A., Durand, T., He, J., Sigal, L., Mori, G.: A variational auto-encoder model for stochastic point processes. In: CVPR (2019)

- 16 Chang et al.
- Miech, A., Zhukov, D., Alayrac, J.B., Tapaswi, M., Laptev, I., Sivic, J.: Howto100m: Learning a text-video embedding by watching hundred million narrated video clips. In: ICCV (2019)
- 24. Oh, J., Guo, X., Lee, H., Lewis, R.L., Singh, S.: Action-conditional video prediction using deep networks in atari games. In: NeurIPS (2015)
- 25. Pathak, D., Agrawal, P., Efros, A.A., Darrell, T.: Curiosity-driven exploration by self-supervised prediction. In: CVPR Workshops (2017)
- Ranzato, M., Szlam, A., Bruna, J., Mathieu, M., Collobert, R., Chopra, S.: Video (language) modeling: a baseline for generative models of natural videos. arXiv preprint arXiv:1412.6604 (2014)
- 27. Rhinehart, N., Kitani, K.M.: First-person activity forecasting with online inverse reinforcement learning. In: ICCV (2017)
- Richard, A., Kuehne, H., Iqbal, A., Gall, J.: Neuralnetwork-viterbi: A framework for weakly supervised video learning. In: CVPR (2018)
- Sener, F., Yao, A.: Zero-shot anticipation for instructional activities. In: ICCV (2019)
- Sermanet, P., Lynch, C., Chebotar, Y., Hsu, J., Jang, E., Schaal, S., Levine, S., Brain, G.: Time-contrastive networks: Self-supervised learning from video. In: ICRA (2018)
- Srinivas, A., Jabri, A., Abbeel, P., Levine, S., Finn, C.: Universal planning networks. In: ICML (2018)
- Sun, C., Myers, A., Vondrick, C., Murphy, K., Schmid, C.: Videobert: A joint model for video and language representation learning. In: ICCV (2019)
- Tang, Y., Ding, D., Rao, Y., Zheng, Y., Zhang, D., Zhao, L., Lu, J., Zhou, J.: Coin: A large-scale dataset for comprehensive instructional video analysis. In: CVPR (2019)
- 34. Vondrick, C., Pirsiavash, H., Torralba, A.: Anticipating visual representations from unlabeled video. In: CVPR (2016)
- 35. Wang, X., Hu, J.F., Lai, J.H., Zhang, J., Zheng, W.S.: Progressive teacher-student learning for early action prediction. In: CVPR (2019)
- Zeng, K.H., Shen, W.B., Huang, D.A., Sun, M., Carlos Niebles, J.: Visual forecasting by imitating dynamics in natural sequences. In: ICCV (2017)
- 37. Zhou, L., Xu, C., Corso, J.J.: Towards automatic learning of procedures from web instructional videos. In: AAAI (2018)
- Zhukov, D., Alayrac, J.B., Cinbis, R.G., Fouhey, D., Laptev, I., Sivic, J.: Cross-task weakly supervised learning from instructional videos. In: CVPR (2019)