

# Learning Procedural-Aware Video Representations Through State-Grounded Hierarchy Unfolding

Jinghan Zhao<sup>1</sup>, Yifei Huang<sup>2</sup>, Feng Lu<sup>1\*</sup>

<sup>1</sup>State Key Laboratory of VR Technology and Systems, School of CSE, Beihang University

<sup>2</sup>The University of Tokyo, Japan

{zhaojinghan,lufeng}@buaa.edu.cn, hyf@iis.u-tokyo.ac.jp

## Abstract

Learning procedural-aware video representations is a key step towards building agents that can reason about and execute complex tasks. Existing methods typically address this problem by aligning visual content with textual descriptions at the task and step levels to inject procedural semantics into video representations. However, due to their high level of abstraction, ‘task’ and ‘step’ descriptions fail to form a robust alignment with the concrete, observable details in visual data. To address this, we introduce ‘states’, *i.e.*, textual snapshots of object configurations, as a visually-grounded semantic layer that anchors abstract procedures to what a model can actually see. We formalize this insight in a novel Task-Step-State (TSS) framework, where tasks are achieved via steps that drive transitions between observable states. To enforce this structure, we propose a progressive pre-training strategy that unfolds the TSS hierarchy, forcing the model to ground representations in states while associating them with steps and high-level tasks. Extensive experiments on the COIN and CrossTask datasets show that our method outperforms baseline models on multiple downstream tasks, including task recognition, step recognition, and next step prediction. Ablation studies show that introducing state supervision is a key driver of performance gains across all tasks. Additionally, our progressive pretraining strategy proves more effective than standard joint training, as it better enforces the intended hierarchical structure.

**Code** — <https://github.com/zhao-jinghan/TSS-unfolding>

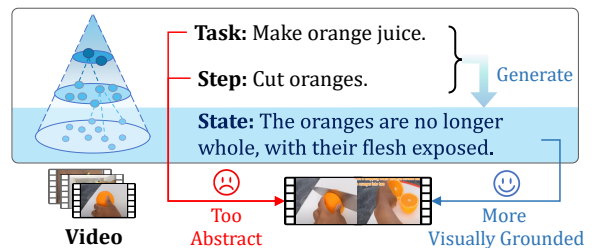
## Introduction

Understanding and performing goal-oriented procedural activities is a core capability for intelligent agents. Humans effortlessly acquire such procedural knowledge from instructional videos, learning not only the overall goals but also the sequence of steps required to achieve them. Enabling machines to learn similar procedural-aware video representations is therefore central to progress in areas such as embodied AI, intelligent assistants, and emerging paradigms like cobodied intelligence (human-AI co-embodied intelligence) (Lu and Zhao 2026).

\*Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

### (a) State-grounded procedural hierarchy



### (b) Learning by hierarchy unfolding

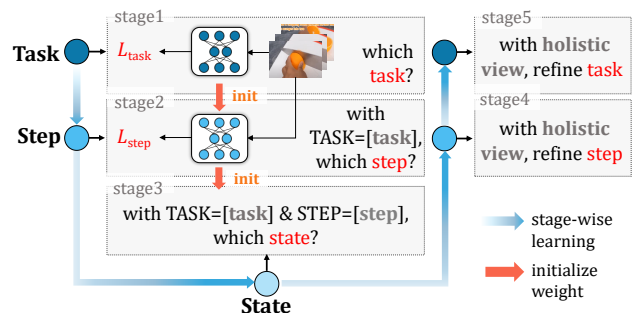


Figure 1: (a) Our method extends the conventional task-step hierarchy by adding a state layer, forming the Task-Step-State (TSS) framework, where state can anchor abstract procedures to visual evidence. (b) To unfold the structure in TSS, we introduce a progressive pre-training strategy by cycling through the order: task→step→state→step→task.

To this end, existing methods typically seek to inject procedural knowledge into video representations by aligning visual content with textual descriptions from external knowledge bases (Lin et al. 2022; Zhou et al. 2023; Mavroudi, Afouras, and Torresani 2023). This alignment is generally performed at two levels of abstraction: the overall task (*e.g.*, “Make orange juice.”) and the constituent steps (*e.g.*, “cut oranges”). By matching video features or ASR transcripts to these descriptions, models learn to recognize the high-level structure of an activity.

However, we argue that this two-level approach suffers from a critical limitation. Due to their high level of abstraction, “task” and “step” descriptions often fail to form a ro-

bust alignment with the concrete, observable details in visual data. A significant semantic gap (Wang et al. 2023a) exists between an abstract command like “cut oranges” and the raw pixels depicting the action. Thus, when training with visual-text alignment, this gap makes it difficult for models to ground abstract procedures in what they can actually see.

To address this, we introduce “states”, *i.e.*, textual snapshots of object configurations, as a visually-grounded semantic layer that anchors abstract procedures to visual evidence. A state describes the observable configuration of key objects, such as “the oranges are no longer whole, with their flesh exposed.”

From a logical standpoint, states form the skeleton of any procedural task. A high-level “Task” can be viewed as a macro-level transition from an initial state to a final state, while a series of “Steps” are the actions that drive the transitions between these intermediate states. Building on this insight, we formalize this structure in a novel Task-Step-State (TSS) framework. In this hierarchy, tasks are achieved via steps, and steps are semantically anchored by the observable state transitions they cause.

To enforce this structure and ensure the model learns the intended hierarchy, we propose a novel progressive pre-training strategy that unfolds the TSS hierarchy. Instead of learning all semantic levels jointly, our strategy forces the model to ground its representations in the most concrete layer—the states. By this means, it better learns to associate states with steps, and finally, steps with high-level tasks. This progressive approach ensures that the learning of abstract concepts is firmly built upon a foundation of visually-grounded understanding.

We validate our approach through extensive experiments on the COIN (Tang et al. 2019) and CrossTask (Zhukov et al. 2019) datasets. Our method outperforms baseline models on multiple downstream tasks, including task recognition, step recognition, and next step prediction. Furthermore, detailed ablation studies confirm our core hypotheses: introducing state supervision is the key driver of these performance gains, and our progressive pre-training strategy is significantly more effective than standard joint training at enforcing the hierarchical structure. Our contributions are:

- We propose a novel Task-Step-State (TSS) framework that introduces “states” as a visually-grounded semantic layer, effectively bridging the semantic gap between abstract procedural texts and concrete visual content.
- We introduce a progressive pre-training strategy that unfolds the TSS hierarchy, ensuring that representations are grounded in concrete states before learning abstract steps and tasks, better enforcing the procedural logic.
- Experiments demonstrate that our method achieves consistent performance gains over existing models and validate that both the TSS framework and the progressive training strategy are key to this success.

## Related Works

### Procedural Learning from Instructional Video

Instructional videos serve as a rich source for learning how agents perceive and act in the world, fueling applications

such as procedure planning (Wang et al. 2023b; Yang, Zhao, and Liu 2025; Ilaslan et al. 2025; Wu et al. 2024), step recognition (Kazakos et al. 2021; Piergiovanni et al. 2021; Zhukov et al. 2019), localization (Dvornik et al. 2022), and task recognition (Ghoddosian, Sayed, and Athitsos 2022). The introduction of large-scale, richly annotated datasets (Zhou, Xu, and Corso 2018; Tang et al. 2019; Afouras et al. 2023) has greatly advanced this field. However, acquiring dense temporal annotations for task-step structures remains a key challenge, limiting the scalability of supervised approaches.

To address this bottleneck, recent efforts have embraced weakly-supervised paradigms that leverage external knowledge sources, such as wikiHow (Koupaee and Wang 2018), to generate supervisory signals for large-scale, unannotated video collections (Lin et al. 2022; Zhou et al. 2023; Mavroudi, Afouras, and Torresani 2023; Samel, Sontakke, and Essa 2025; Zhong et al. 2023). These approaches typically aim to learn procedural-aware video representations by aligning visual or transcribed-text features with textual descriptions of steps. For instance, Paprika (Zhou et al. 2023) utilizes pre-trained visual and textual encoders to perform cross-modal matching, and builds a graph that captures the temporal progression of both visual content and procedural steps derived from wikiHow. Similarly, (Mavroudi, Afouras, and Torresani 2023) integrates video frames, ASR-generated narrations, and step descriptions, proposing dual alignment paths (“step-frame” and “step-narration-frame”) to bridge the gap between vision and language modalities.

While effective at capturing the high-level procedure, these approaches operate purely at the abstract task and step levels. This introduces a critical limitation, as noted by works like (Wang et al. 2023a): a semantic gap exists between abstract, action-centric text and concrete visual data. Our work directly confronts this gap. We argue that to truly ground procedural learning, the knowledge hierarchy itself must be augmented. We achieve this by introducing “states” as a new, visually-grounded semantic layer, bridging the gap between abstract commands and observable visual content.

### State-Based/Object-Centric Video Understanding

The concept of modeling object states and their transitions is a cornerstone of robust action understanding (Aboubakr, Crowley, and Ronfard 2019; Saini et al. 2022; Doughty et al. 2020), procedural planning (Wang et al. 2023b; Zhao et al. 2022; Ilaslan et al. 2025; Wu et al. 2024), and beyond. Recent work has focused on automatically extracting and utilizing state information from video (Alayrac et al. 2017; Souček et al. 2022a,b; Xue, Ashutosh, and Grauman 2024; Niu et al. 2024). For instance, (Souček et al. 2024) mines triplets of (initial state, action, end state) from instructional videos to train generative models of object states. (Niu et al. 2024) represent each step as a change of state in procedural planning, supplement state descriptions via Large Language Model (LLM), and align them with visual observations to achieve tracking of state changes.

Our work builds directly on these insights. We adopt the effective strategy of using LLMs to generate state descriptions, but move beyond simply tracking them. We are the first to formalize states as a fundamental component within a

three-tiered hierarchy (Task-Step-State), specifically for pre-training procedural-aware video representations.

### Progressive and Curriculum-Based Learning

Progressive learning draws inspiration from the way humans acquire skills, where knowledge is built up gradually in a meaningful order. In machine learning, this idea has been formalized primarily through two paradigms: curriculum learning, which trains models from easy to hard examples (Bengio et al. 2009; Hachohen and Weinshall 2019), and hard example mining, which does the reverse by focusing on difficult samples first (Sun and Pang 2018; Suh et al. 2019). In recent years, progressive learning has been applied in various domains from image generation to gait recognition (Du et al. 2021; Wang et al. 2018; Dou et al. 2022; Garg, Kumar, and Rawat 2025). Recent work in video understanding has also adopted this principle. For instance, STPro (Garg, Kumar, and Rawat 2025) uses a spatio-temporal curriculum, progressively increasing the complexity of action sequences and scene backgrounds to improve grounding.

Our work introduces a novel application of this progressive philosophy. Our proposed pre-training pathway enables the model to unfold structured procedural knowledge within our TSS framework, leading to more robust and generalizable procedural-aware video representations.

### Method

Our method is designed around two core contributions to learn procedural-aware video representations. We first introduce the Task-Step-State (TSS) framework, a new knowledge structure that grounds abstract procedures in observable visual content. We then propose a novel progressive pre-training strategy that teaches a model to understand this hierarchy by unfolding it in a principled, stage-wise manner.

### Problem Formulation

We aim to learn robust video representations guided by structured procedural knowledge. Starting from a textual corpus, such as wikiHow (Koupaee and Wang 2018), consisting of tasks decomposed into sequential steps, our goal is to train a video encoder  $\Phi$  mapping video clips into representations useful for procedural understanding tasks (e.g., task recognition, step recognition, and forecasting).

Formally, given video clips  $v$ , the encoder  $\Phi$  produces representations:  $f_v = \Phi(v)$ . Representation quality is assessed through downstream task performance.

### TSS Framework Construction

The foundation of our approach is the creation of a new, state-augmented knowledge framework to bridge the semantic gap between abstract steps and concrete visual content.

**Grounding Procedures in States** The conventional structure for procedural knowledge, used in prior work, is a two-level hierarchy. A high-level *task* (denoted by  $T$ ) is decomposed into an ordered sequence of *steps* (denoted by  $s$ ). However, as we have established, the purely abstract nature of these descriptions is the source of the critical semantic

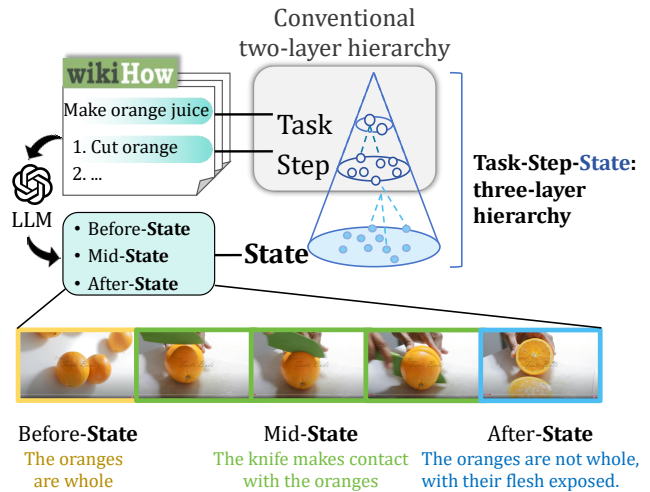


Figure 2: The TSS framework extends the task-step hierarchy with an LLM-generated state layer.

gap that hampers robust alignment with visual data. To resolve this, we augment the structure by introducing a third, more visually grounded layer: the *state*, denoted by  $c$ . We define a state as a *textual snapshot of object configurations*. This includes not only the properties of single objects, such as their attributes, quantity, and shape, but also the spatial relationships between them. To capture the temporal nature of actions, denote the  $j$ -th step of the  $i$ -th task as  $s_{i,j}$ , we associate each  $s_{i,j}$  from the original corpus with three distinct state types: a before-state ( $c_{i,j}^b$ ), a mid-state ( $c_{i,j}^m$ ), and an after-state ( $c_{i,j}^a$ ). This enriches each step into an expanded tuple:

$$s'_{i,j} = (s_{i,j}, c_{i,j}^b, c_{i,j}^m, c_{i,j}^a),$$

forming our proprietary three-level hierarchy.

Specifically, we construct our state-augmented knowledge base from the wikiHow corpus. As manual annotation of states for every step is infeasible, we employ a Large Language Model (LLM) GPT-4o-mini to automatically infer and generate the state descriptions, which is a practical approach with precedents in related work (Niu et al. 2024). For each step description  $s_{i,j}$ , the LLM is prompted to generate the corresponding before-state, mid-state, and after-state. To ensure the logical consistency and quality of the generated text, we utilize a Chain-of-Thought (CoT) (Wei et al. 2022) prompting strategy. This automated process effectively transforms the standard wikiHow corpus into our rich, state-aware TSS framework, which we will then use for the next step: pseudo-label generation.

**Pseudo-Label Generation from TSS** To generate supervisory signals for our pre-training, we align a large corpus of unannotated videos with our Task-Step-State (TSS) framework. This process creates a rich set of matched video segments and the textual descriptions of tasks, steps, and states, which can be further used in pre-training.

We begin by projecting video data and text descriptions into a shared embedding space. Using off-the-shelf pre-

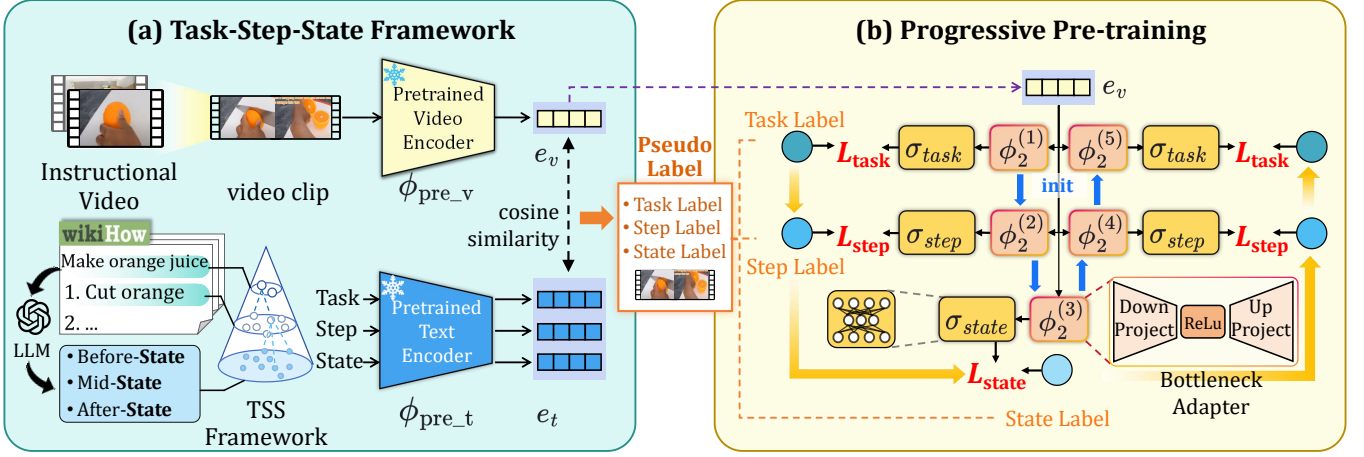


Figure 3: Overview. (a) First, an LLM generates state descriptions from the task and step in WikiHow, yielding the three-level TSS framework. We encode these texts with a frozen text encoder  $\phi_{\text{pre}_t}$  and encode video clips from HowTo100M with a frozen vision encoder  $\phi_{\text{pre}_v}$ , producing text features  $e_t$  and video features  $e_v$ . Cosine similarity between  $e_t$  and  $e_v$  creates pseudo labels at each level. (b) Second, the progressive pre-training strategy trains the vision model in five stages, each focused on one TSS layer. Every stage takes  $e_v$ , passes it through an adapter  $\phi_2^{(i)}$  and a task-specific head  $\sigma$ , and is supervised by the corresponding pseudo label.

trained vision and text encoders  $\phi_{\text{pre}_v}(\cdot)$  and  $\phi_{\text{pre}_t}(\cdot)$ , we obtain initial embeddings  $e_v = \phi_{\text{pre}_v}(v)$  and  $e_t = \phi_{\text{pre}_t}(t)$  for each video clip  $v$  and text description  $t$ . The correspondence between these initial embeddings is measured by their cosine similarity:

$$S(e_v, e_t) = \frac{e_v \cdot e_t}{|e_v| |e_t|} \quad (1)$$

To create a stable and canonical set of targets, we then cluster the text embeddings  $e_t$  at each semantic level. For instance, all step description embeddings are clustered to form a set of representative ‘‘step nodes’’,  $\mathcal{N}_{\text{step}}$ . This process is repeated for tasks ( $\mathcal{N}_{\text{task}}$ ) and for each type of the before/mid/after state ( $\mathcal{N}_{\text{state}}$ ) separately, to yield  $\mathcal{N}_{\text{state}}^b$ ,  $\mathcal{N}_{\text{state}}^m$  and  $\mathcal{N}_{\text{state}}^a$ .

With these semantic nodes established, we generate pseudo-labels by aligning each video clip embedding  $e_v$  to them. We calculate the similarity of  $e_v$  to all nodes within a given set (e.g.,  $\mathcal{N}_{\text{step}}$ ) and select the top-k most similar nodes as positive targets for a multi-class classification objective. This yields our primary supervision labels of Video-Node Matching:

- **taskVNM**: Generated by matching video features against task nodes ( $\mathcal{N}_{\text{task}}$ ).
- **stepVNM**: Generated by matching video features against step nodes ( $\mathcal{N}_{\text{step}}$ ).
- **stateVNM**: First, the node set ( $\mathcal{N}_{\text{state}}^b / \mathcal{N}_{\text{state}}^m / \mathcal{N}_{\text{state}}^a$ ) is determined by the best-matching node, and then the top-k matching nodes are retrieved.

Finally, inspired by prior work on learning procedural graphs (Zhou et al. 2023), we extract two additional supervision signals to capture richer step-level relationships. By analyzing the graph structure formed from step nodes, we derive:

- **stepNRL** (Node Relation Learning) : trains the model to identify a step’s preceding and succeeding nodes.
- **stepTCL** (Task Context Learning) : trains the model to recognize other step nodes belonging to the same overarching task.

### Progressive Pre-Training

One critical aspect of our method lies in how we teach the model the TSS hierarchy. Instead of a black-box joint training approach, we propose a principled, progressive pre-training strategy that unfolds the hierarchy in a specific, highly effective pathway.

**Model Architecture and Preliminaries** Our model, denoted as  $\Phi$ , uses a parameter-efficient design to make large-scale pretraining computationally affordable. It is composed of two key components:

- **A frozen pre-trained vision encoder,  $\phi_v$ .** In our work, we use vision encoder of S3D model (Miech et al. 2020) simultaneously as  $\phi_v$  and  $\phi_{\text{pre}_v}$ , keeping its weights frozen throughout our training.
- **A trainable lightweight adapter  $\phi_2$ .** The adapter features a bottleneck architecture (Houlsby et al. 2019; Sung, Cho, and Bansal 2022), consisting of a down-projection layer followed by an up-projection layer. This lightweight module is the primary component we train, allowing it to adapt the generic visual features from  $\phi_v$  for our specific procedural understanding tasks.

The final video representation  $f_v$  is thus computed as:

$$f_v = \Phi(v) = \phi_2(\phi_v(v)) = \phi_2(e_v)$$

This representation is then fed into various **task-specific heads**, denoted by  $\sigma(\cdot)$ , which are simple Multi-Layer Perceptrons (MLPs) designed for multi-class classification objectives defined by our pseudo-labels.

**Progressive Pre-training Strategy** With this architecture in place, we address a fundamental question: *how can the intrinsic, hierarchical connections between the Task, Step, and State layers be best exploited for representation learning?* A standard approach of jointly training on all supervisory signals at once treats the learning process as a black box, offering little insight into how these distinct knowledge types influence one another.

To answer this question, we move beyond joint training and propose a **progressive pre-training methodology**. Our work extends this concept to explore the optimal learning pathway for a hierarchical knowledge system like TSS. Our work asks more complex questions: Does knowledge transfer exclusively downward (e.g., from Task to Step), or can knowledge from a lower level (like State) be used to retrospectively refine an understanding of a higher level (like Step or even Task)?

To systematically explore these questions, we define a training process that proceeds in a sequence of stages. This approach allows us to control the flow of information and investigate the impact of different learning pathways. The mechanism is as follows:

- **Staged Learning:** The overall training is decomposed into a series of stages. Each stage  $i$  is designed to learn a specific procedural knowledge, and corresponds to different task-specific heads ( $\sigma_{task}/\sigma_{step}/\sigma_{state}$ ).
- **Knowledge Propagation via Adapter:** In each stage, we train the adapter ( $\phi_2^{(i)}$ ) alongside a set of randomly initialized task heads required for that stage’s objectives. Upon completion, we discard the task heads but preserve the weights of the trained adapter to initialize the adapter ( $\phi_2^{(i+1)}$ ) for the subsequent stage ( $i+1$ ). The task heads for the new stage are again initialized randomly.

This progressive process, where knowledge is accumulated and passed through the adapter, allows us to assess different learning trajectories.

Through systematic exploration, we design our primary learning strategy around a full cycle: **Task**→**Step**→**State**→**Step**→**Task**. This pathway is designed to model a complete analysis-synthesis process. It first performs top-down analysis, using high-level context to predict fine-grained details, before performing bottom-up synthesis, using grounded details to verify and confirm high-level hypotheses. This circular pathway enables bidirectional information flow, for instance, allowing state-level knowledge to refine step representations. Furthermore, each segment of this cycle can be treated as an independent progressive strategy, enabling us to ablate and evaluate the contribution of each knowledge transfer direction in our experiments.

## Experiments

We conduct a comprehensive set of experiments designed to first dissect our proposed method and then validate its superiority. We begin with detailed ablation studies to answer two fundamental questions: **(1)** What is the optimal progressive pathway for learning the Task-Step-State hierarchy? **(2)** How does each semantic layer contribute to the final repre-

sentation? Having identified the best-performing configuration from these studies, we then **(3)** compare our final model against state-of-the-art methods to demonstrate its effectiveness on standard procedural understanding benchmarks.

## Datasets

For pre-training, our visual data is an 85K video subset of HowTo100M (Miech et al. 2019), following (Bertasius, Wang, and Torresani 2021). Our textual data is a curated subset of WikiHow (Koupaee and Wang 2018) comprising 1,053 tasks and 10,588 steps. We evaluate our representations on two standard benchmarks: COIN (Tang et al. 2019) and CrossTask (Zhukov et al. 2019). COIN consists of 11.8K videos across 180 tasks. CrossTask includes 83 tasks, of which only 18 have full temporal annotations.

## Evaluation Protocol

Consistent with prior works (Zhou et al. 2023; Lin et al. 2022), we assess the quality of our learned representation on three downstream tasks: **Task Recognition (TR)**, **Step Recognition (SR)**, and **Step Forecasting (SF)**. We report Top-1 Accuracy as the primary metric. For fine-tuning, we attach two distinct evaluation models to our frozen representation: a simple MLP and a more powerful one-layer Transformer (Vaswani et al. 2017). This allows us to assess the robustness of our learned features across downstream models of varying capacities. The specific architectural details for both evaluation models are deferred to the supplementary material in our GitHub repository.

## Implemental Details

For data processing, we segment the source videos into non-overlapping 9.6-second clips. For each of the 10,588 procedural steps, we generate three types of state descriptions, resulting in a total of 31,764 state-level texts. The video feature  $e_v$  and text feature  $e_t$  both have a dimensionality of 512. Our adapter employs a bottleneck architecture, projecting the 512-dimensional features down to a hidden dimension of 128 before restoring them to 512. We train our model with a batch size of 256 using the Adam optimizer. All experiments were conducted on a single server equipped with eight H200 GPUs. Further details regarding hyperparameters and training procedures are provided in our GitHub repository.

## Pathway and Component Analysis

We first conduct an extensive analysis to understand the impact of our core design choices, focusing on the pre-training curriculum.

ID	Pre-training Pathway
Path-1	Task
Path-2	Task→Step
Path-3 (Ours)	Task→Step→State
Path-4 (Ours)	Task→Step→State→Task
Path-5 (Ours)	Task→Step→State→Step
Path-6 (Ours)	Task→Step→State→Step→Task

Table 1: Proposed progressive pre-training pathways

Pretrain Pathway	Downstream MLP						Downstream Transformer					
	COIN			CrossTask			COIN			CrossTask		
	TR	SR	SF	TR	SR	SF	TR	SR	SF	TR	SR	SF
No pretrain	2.09	1.37	0.84	53.00	20.21	23.27	78.31	39.23	35.43	89.44	56.82	56.17
Path-1	73.31	34.18	23.67	89.44	52.09	52.40	82.25	37.84	37.70	89.03	54.58	53.67
Path-2	82.45	43.06	36.04	89.44	56.55	56.37	83.30	40.96	40.03	90.06	55.71	55.83
Path-3 (Ours)	80.73	41.28	34.35	89.65	55.16	55.27	82.16	40.24	39.22	88.82	54.84	54.13
Path-4 (Ours)	77.74	37.51	24.84	89.86	54.29	54.57	81.16	36.04	36.85	89.65	55.34	54.20
Path-5 (Ours)	83.78	44.54	38.07	89.44	57.92	57.13	83.11	42.42	40.40	89.44	57.08	56.50
Path-6 (Ours)	83.30	44.04	36.94	90.27	57.63	57.13	83.25	42.52	40.02	89.44	57.08	56.50
Mix-Train	77.74	38.43	29.79	89.23	53.02	53.17	82.92	37.93	38.13	89.03	54.06	53.70
Fusion-Concat	83.16	44.21	36.65	89.65	57.34	57.27	82.49	42.34	40.10	90.27	56.70	56.17
Fusion-AvgPool	83.11	44.35	36.88	90.06	57.48	57.50	83.40	42.50	40.20	90.68	56.84	56.10

Table 2: Ablation results. (1) We compare the downstream accuracy of six progressive pre-training pathways and joint training. (2) We fuse video representations pre-trained with three different levels and evaluate two fusion strategies.

**Analysis of the Pre-training Pathway** To investigate the optimal pre-training curriculum for incorporating hierarchical procedural knowledge, we first experiment with six distinct pre-training paths and compare their performance on downstream tasks. The sequential paths investigated are listed in Table 1.

The results are presented in Table 2. Additionally, we report the results of MIL-NCE (Miech et al. 2020) to serve as our “No pretrain” baseline for comparison. MIL-NCE is a pre-training objective for video-narration contrastive learning. For this baseline, we employ the frozen S3D model provided by the authors to extract video features. We visualized the performance matrix for all ablation settings as a heatmap (Figure 4) for better visibility.

From Table 2, it is clear that pre-training on task and step information is beneficial, with Path-2 outperforming Path-1 and the “No Pre-train” baseline. This aligns with prior work. More insights emerge when we introduce the state layer.

Path-3 (Task  $\rightarrow$  Step  $\rightarrow$  State) is our most direct state-aware curriculum. This pathway reveals a crucial challenge: while grounding the representation in states is the goal, doing so naively causes a slight performance dip on the more abstract task and step recognition tasks. This happens because the model’s focus shifts heavily towards the concrete vision-state alignment, temporarily de-emphasizing the higher-level semantics. This finding validates our core premise: simply adding states is not enough; how they are integrated is critical.

This is further evidenced by Path-4 (...  $\rightarrow$  State  $\rightarrow$  Task). Transitioning directly from the concrete state layer back to the highly abstract task layer results in comparable or even lower performance compared to Path-3 across most settings (e.g., TR of 81.16 on COIN). This observation powerfully supports the architectural rationale of our three-layer hierarchy, confirming that a significant semantic gap exists between states and tasks, and that the step layer serves as an essential intermediate bridge.

Our U-shaped pathway is designed to correctly utilize this bridge. Path-5 (...  $\rightarrow$  State  $\rightarrow$  Step) demonstrates the power of this analysis-synthesis approach. By using the now-grounded state knowledge to retrospectively refine its

understanding of step semantics, it achieves top-tier performance, with Step Recognition (SR) and Step Forecasting (SF) scores of 42.42 and 40.40 on COIN, respectively. This result reveals two key insights:

- First, it confirms that state-level pre-training forces the model to anchor its understanding in observable evidence, which serves as a robust foundation for interpreting more abstract concepts.
- Second, it validates that our progressive pre-training is effective for transferring this visual grounding upward to enrich and refine the abstract step representations.

Path-6 extends Path-5 by incorporating the final Step  $\rightarrow$  Task traceback. This addition only leads to marginal performance gains. We hypothesize that this is due to convergence, as the step representations may have already been sufficiently optimized in Path-5, thus limiting the marginal gain from the final, high-level objective.

**Analysis of Alternative Learning Strategies** Having established that our TSS framework progressively injects procedure knowledge to strengthen video representations, we next explore: can the same knowledge be fully exploited (1) without the staged curriculum, or (2) by directly fusing representations learned at different semantic levels? To answer this, we conduct two ablation studies in Table 2:

- Joint training: we investigate an alternative training strategy: joint pre-training (referred to as “**Mix-Train**”). In this setup, the model is supervised simultaneously by all objectives from the task, step, and state levels, rather than progressively.
- Multi-Level Feature Fusion: We try to integrate features learned from different stages of our progressive pre-training approach. Specifically, we investigate whether the three video representations obtained through distinct pre-training pathways ( $f_v^{task}$ ,  $f_v^{step}$ , and  $f_v^{state}$ , derived from Path-1, Path-2, and Path-3, respectively) can be fused to achieve improved performance. We evaluated two fusion methods (Figure 5): (1) concatenation along the temporal dimension, and (2) average pooling.

Pretrain Method	Downstream MLP						Downstream Transformer					
	COIN			CrossTask			COIN			CrossTask		
	TR	SR	SF	TR	SR	SF	TR	SR	SF	TR	SR	SF
MIL-NCE (Miech et al. 2020)	2.09	1.37	0.84	53.00	20.21	23.27	78.31	39.23	35.43	89.44	56.82	56.17
Paprika (Zhou et al. 2023)	81.54	42.39	34.10	89.65	56.21	55.77	82.83	41.19	38.93	90.27	55.57	55.67
<b>Ours</b> Path-5	<b>83.78</b>	<b>44.54</b>	<b>38.07</b>	89.44	<b>57.92</b>	<b>57.13</b>	<b>83.11</b>	<b>42.42</b>	<b>40.40</b>	89.44	<b>57.08</b>	<b>56.50</b>
Gains to Paprika	<u>+2.24</u>	<u>+2.15</u>	<u>+3.97</u>	<u>-0.21</u>	<u>+1.71</u>	<u>+1.36</u>	<u>+0.28</u>	<u>+1.23</u>	<u>+1.47</u>	<u>-0.83</u>	<u>+1.51</u>	<u>+0.83</u>
Path-6	<b>83.30</b>	<b>44.04</b>	<b>36.94</b>	<b>90.27</b>	<b>57.63</b>	<b>57.13</b>	<b>83.25</b>	<b>42.52</b>	<b>40.02</b>	89.44	<b>57.08</b>	<b>56.50</b>
Gains to Paprika	<u>+1.76</u>	<u>+1.65</u>	<u>+2.84</u>	<u>+0.62</u>	<u>+1.42</u>	<u>+1.36</u>	<u>+0.42</u>	<u>+1.33</u>	<u>+1.09</u>	<u>-0.83</u>	<u>+1.51</u>	<u>+0.83</u>

Table 3: Accuracy of our method and the SOTA baselines on downstream procedural understanding tasks. Our method surpasses Paprika on all three tasks.

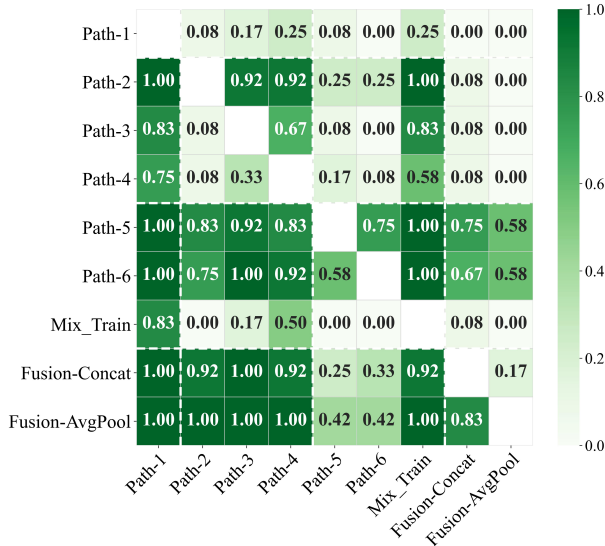


Figure 4: Ablation results Comparison. The value at position  $(i, j)$  in the heat map denotes the fraction of 12 evaluation settings where configuration  $i$  outperforms configuration  $j$ .

For Mix\_Train strategy, its accuracy falls between Path-1 and Path-2. This result demonstrates that simply training on all objectives altogether is not optimal. This suggests that the progressive nature of our pathway is crucial, as it allows the model to effectively leverage the structural and causal relationships between the three semantic levels. The direct joint training approach fails to capture these relationships.

Regarding the two fusion approaches, the Fusion-AvgPool model is very strong. This strong performance provides further validation for our central thesis, demonstrating that the state layer contributes critical, complementary information that is missing from the standard Task-Step hierarchy. Meanwhile, our progressive model, Path-5, is still superior, particularly with downstream MLP. This proves that while adding state information is crucial, our progressive pre-training strategy is a more effective method for integrating this knowledge into a single, cohesive representation.

**Comparison to the State of the Art** Having validated our TSS framework and identified Path-5 and 6 as our champion models, we now compare it against state-of-the-art (SOTA)




Figure 5: Different ways to fuse the three visual representations after adapter fine-tuning in the ablation study.

baselines. As shown in Table 3, for a fair comparison, all methods share the same frozen S3D encoder and downstream model architectures. Our primary SOTA baseline is Paprika (Zhou et al. 2023), which also leverages wikiHow but operates only at the step level.

Results are reported in Table 3, where our method consistently outperforms the state of the art. On the COIN dataset with a Transformer head, our Path-5 model surpasses Paprika by +0.28, +1.23, and +1.47 points on Task Recognition, Step Recognition, and Step Forecasting, respectively. Our slightly larger Path-6 model extends these gains even further. The improvements are even more pronounced with a simpler MLP head, where our method achieves gains of up to +2.24 points over Paprika. This confirms that representations pre-trained with our state-grounded, progressive method encode more salient and robust procedural knowledge, validating our overall approach.

## Conclusion

In this work, we address a fundamental limitation in procedural video understanding: the failure of abstract task and step descriptions to align with concrete visual data. We introduce the TSS framework, a novel hierarchy that resolves this by grounding abstract procedures in observable, visually salient states. To teach this structure, we develop a progressive pre-training strategy that effectively unfolds the hierarchy, forcing the model to first ground its understanding in states before synthesizing knowledge of steps and tasks. Our state-of-the-art results on multiple benchmarks validate our central thesis: grounding abstract procedures in concrete, observable states is a crucial and previously missing component for robust procedural representation learning.

## Acknowledgements

This work was supported by Beijing Natural Science Foundation (L242019) and JSPS KAKENHI JP25K24384.

## References

- Aboubakr, N.; Crowley, J. L.; and Ronfard, R. 2019. Recognizing manipulation actions from state-transformations. *arXiv preprint arXiv:1906.05147*.
- Afouras, T.; Mavroudi, E.; Nagarajan, T.; Wang, H.; and Torresani, L. 2023. Ht-step: Aligning instructional articles with how-to videos. *Advances in Neural Information Processing Systems*, 36: 50310–50326.
- Alayrac, J.-B.; Laptev, I.; Sivic, J.; and Lacoste-Julien, S. 2017. Joint discovery of object states and manipulation actions. In *Proceedings of the IEEE International Conference on Computer Vision*, 2127–2136.
- Bengio, Y.; Louradour, J.; Collobert, R.; and Weston, J. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, 41–48.
- Bertasius, G.; Wang, H.; and Torresani, L. 2021. Is space-time attention all you need for video understanding? In *icml*, volume 2, 4.
- Dou, H.; Zhang, P.; Zhao, Y.; Dong, L.; Qin, Z.; and Li, X. 2022. Gaitmpl: Gait recognition with memory-augmented progressive learning. *IEEE Transactions on Image Processing*, 33: 1464–1475.
- Doughty, H.; Laptev, I.; Mayol-Cuevas, W.; and Damen, D. 2020. Action modifiers: Learning from adverbs in instructional videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 868–878.
- Du, R.; Xie, J.; Ma, Z.; Chang, D.; Song, Y.-Z.; and Guo, J. 2021. Progressive learning of category-consistent multi-granularity features for fine-grained visual classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12): 9521–9535.
- Dvornik, N.; Hadji, I.; Pham, H.; Bhatt, D.; Martinez, B.; Fazly, A.; and Jepson, A. D. 2022. Graph2vid: Flow graph to video grounding for weakly-supervised multi-step localization. *arXiv preprint arXiv:2210.04996*.
- Garg, A.; Kumar, A.; and Rawat, Y. S. 2025. Stpro: Spatial and temporal progressive learning for weakly supervised spatio-temporal grounding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 3384–3394.
- Ghoddosian, R.; Sayed, S.; and Athitsos, V. 2022. Hierarchical modeling for task recognition and action segmentation in weakly-labeled instructional videos. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 1922–1932.
- Hacohen, G.; and Weinshall, D. 2019. On the power of curriculum learning in training deep networks. In *International conference on machine learning*, 2535–2544. PMLR.
- Houlsby, N.; Giurgiu, A.; Jastrzebski, S.; Morrone, B.; De Laroussilhe, Q.; Gesmundo, A.; Attariyan, M.; and Gelly, S. 2019. Parameter-efficient transfer learning for NLP. In *International conference on machine learning*, 2790–2799. PMLR.
- Haslan, M. F.; Köksal, A.; Lin, K. Q.; Satar, B.; Shou, M. Z.; and Xu, Q. 2025. VG-TVP: Multimodal Procedural Planning via Visually Grounded Text-Video Prompting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 3886–3894.
- Kazakos, E.; Huh, J.; Nagrani, A.; Zisserman, A.; and Damen, D. 2021. With a little help from my temporal context: Multimodal egocentric action recognition. *arXiv preprint arXiv:2111.01024*.
- Koupaee, M.; and Wang, W. Y. 2018. Wikihow: A large scale text summarization dataset. *arXiv preprint arXiv:1810.09305*.
- Lin, X.; Petroni, F.; Bertasius, G.; Rohrbach, M.; Chang, S.-F.; and Torresani, L. 2022. Learning to recognize procedural activities with distant supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13853–13863.
- Lu, F.; and Zhao, Q. 2026. Towards embodied/symbolic AI: concept and eight scientific and technical problems. *SCIENCE CHINA Information Sciences*, 69(1): 116101:1–116101:3.
- Mavroudi, E.; Afouras, T.; and Torresani, L. 2023. Learning to ground instructional articles in videos through narrations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15201–15213.
- Miech, A.; Alayrac, J.-B.; Smaira, L.; Laptev, I.; Sivic, J.; and Zisserman, A. 2020. End-to-end learning of visual representations from uncurated instructional videos. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9879–9889.
- Miech, A.; Zhukov, D.; Alayrac, J.-B.; Tapaswi, M.; Laptev, I.; and Sivic, J. 2019. Howto100m: Learning a text-video embedding by watching hundred million narrated video clips. In *Proceedings of the IEEE/CVF international conference on computer vision*, 2630–2640.
- Niu, Y.; Guo, W.; Chen, L.; Lin, X.; and Chang, S.-F. 2024. Schema: State changes matter for procedure planning in instructional videos. *arXiv preprint arXiv:2403.01599*.
- Piergiovanni, A.; Angelova, A.; Ryoo, M. S.; and Essa, I. 2021. Unsupervised discovery of actions in instructional videos. *arXiv preprint arXiv:2106.14733*.
- Saini, N.; He, B.; Shrivastava, G.; Rambhatla, S. S.; and Shrivastava, A. 2022. Recognizing actions using object states. In *ICLR2022 Workshop on the Elements of Reasoning: Objects, Structure and Causality*, volume 1, 3.
- Samel, K.; Sontakke, N.; and Essa, I. 2025. Leveraging procedural knowledge and task hierarchies for efficient instructional video pre-training. *arXiv preprint arXiv:2502.17352*.
- Souček, T.; Alayrac, J.-B.; Miech, A.; Laptev, I.; and Sivic, J. 2022a. Look for the change: Learning object states and state-modifying actions from untrimmed web videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13956–13966.
- Souček, T.; Alayrac, J.-B.; Miech, A.; Laptev, I.; and Sivic, J. 2022b. Multi-task learning of object state changes from uncurated videos. *arXiv preprint arXiv:2211.13500*.

- Souček, T.; Damen, D.; Wray, M.; Laptev, I.; and Sivic, J. 2024. Genhowto: Learning to generate actions and state transformations from instructional videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6561–6571.
- Suh, Y.; Han, B.; Kim, W.; and Lee, K. M. 2019. Stochastic class-based hard example mining for deep metric learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7251–7259.
- Sun, H.; and Pang, Y. 2018. GlanceNets—Efficient convolutional neural networks with adaptive hard example mining. *Science China. Information Sciences*, 61(10): 109101.
- Sung, Y.-L.; Cho, J.; and Bansal, M. 2022. Vi-adapter: Parameter-efficient transfer learning for vision-and-language tasks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5227–5237.
- Tang, Y.; Ding, D.; Rao, Y.; Zheng, Y.; Zhang, D.; Zhao, L.; Lu, J.; and Zhou, J. 2019. Coin: A large-scale dataset for comprehensive instructional video analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1207–1216.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, A.-L.; Lin, K.-Y.; Du, J.-R.; Meng, J.; and Zheng, W.-S. 2023a. Event-guided procedure planning from instructional videos with text supervision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13565–13575.
- Wang, H.; Wu, Y.; Guo, S.; and Wang, L. 2023b. Pdpp: Projected diffusion for procedure planning in instructional videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14836–14845.
- Wang, Y.; Perazzi, F.; McWilliams, B.; Sorkine-Hornung, A.; Sorkine-Hornung, O.; and Schroers, C. 2018. A fully progressive approach to single-image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 864–873.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35: 24824–24837.
- Wu, Y.; Wang, H.; Wang, J.; and Wang, L. 2024. Open-event procedure planning in instructional videos. *arXiv preprint arXiv:2407.05119*.
- Xue, Z.; Ashutosh, K.; and Grauman, K. 2024. Learning object state changes in videos: An open-world perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18493–18503.
- Yang, D.; Zhao, Z.; and Liu, Y. 2025. Planllm: Video procedure planning with refinable large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 9166–9174.
- Zhao, H.; Hadji, I.; Dvornik, N.; Derpanis, K. G.; Wildes, R. P.; and Jepson, A. D. 2022. P3iv: Probabilistic procedure planning from instructional videos with weak supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2938–2948.
- Zhong, Y.; Yu, L.; Bai, Y.; Li, S.; Yan, X.; and Li, Y. 2023. Learning procedure-aware video representation from instructional videos and their narrations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14825–14835.
- Zhou, H.; Martín-Martín, R.; Kapadia, M.; Savarese, S.; and Niebles, J. C. 2023. Procedure-aware pretraining for instructional video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10727–10738.
- Zhou, L.; Xu, C.; and Corso, J. 2018. Towards automatic learning of procedures from web instructional videos. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Zhukov, D.; Alayrac, J.-B.; Cinbis, R. G.; Fouhey, D.; Laptev, I.; and Sivic, J. 2019. Cross-task weakly supervised learning from instructional videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3537–3545.