EgoOops: A Dataset for Mistake Action Detection from Egocentric Videos with Procedural Texts

Yuto Haneji[1], Taichi Nishimura[2], Hirotaka Kameko[1], Keisuke Shirai[1], Tomoya Yoshida[1],
Keiya Kajimura[1], Koki Yamamoto[1], Taiyu Cui[1], Tomohiro Nishimoto[1], Shinsuke Mori[1]
[1]Kyoto University, [2]LY Corporation

haneji.yuto.s66@kyoto-u.jp, tainishi@lycorp.co.jp,
shirai.keisuke.5y@kyoto-u.ac.jp, {kameko,forest}@i.kyoto-u.ac.jp
https://y-haneji.github.io/EgoOops-project-page/

Egocentric video

Procedural text

Task: cardboard craft

1. "Draw auxiliary lines of the shape and the fold on cardboard according to dimensional drawings"
2. "Cut along the auxiliary lines of the shape"
3. "Leave the ends of the liner for gluing."
4. "Make incisions on the auxiliary lines of the fold."
5. "Fold along the incisions of the fold."
6. "Adhere the liner with glue."
7. …

Steps: 3 | 4 | 5

Mistake labels: Unintended action | N/A (correct) | Working in the wrong way

Descriptions: "Cut off the end of the liner." | N/A | "Ignored incisions."

Figure 1. Overview of the EgoOops dataset. EgoOops includes 50 egocentric videos across five procedural domains and corresponding procedural texts. It contains three types of annotations: video-text alignment, mistake labels, and descriptions explaining the errors in each segment.

Abstract

Mistake action detection from egocentric videos is crucial for developing intelligent archives that detect workers' errors and provide feedback. Previous studies have been limited to specific domains, focused on detecting mistakes from videos without procedural texts, and analyzed whether actions are mistakes. To address these limitations, in this paper, we propose the EgoOops dataset, which includes egocentric videos, procedural texts, and three types of annotations: video-text alignment, mistake labels, and descriptions for mistakes. EgoOops covers five procedural domains and includes 50 egocentric videos. The video-text alignment allows the model to detect mistakes based on both videos and procedural texts. The mistake labels and descriptions enable detailed analysis of real-world mistakes. Based on EgoOops, we tackle two tasks: video-text alignment and mistake detection. For video-text alignment, we enhance the recent StepFormer model with an additional loss for fine-tuning. Based on the alignment results, we propose a multimodal classifier to predict mistake labels. In our experiments, the proposed methods achieve higher performance than the baselines. In addition, our ablation study demonstrates the effectiveness of combining videos and texts. We will release the dataset and codes upon publication.

1. Introduction

When performing procedural tasks such as assembly or scientific experiments, we follow procedural texts to carry them out in the real world. During this process, mistakes negatively impact quality, speed, cost, and safety. Common errors include skipping steps or performing incorrect actions, which can sometimes result in life-or-death situations. To reduce them, the goal of this study is to develop an intelligent video archive that records workers' activities,

detects their mistakes, and provides feedback to them.

To record worker's activities, egocentric views gain significant attention because they contain detailed procedural activities to achieve a goal [2, 6, 14, 19, 20, 26, 28, 30, 31, 34, 37, 42]. Recent work collected large-scale egocentric video datasets that consist of procedural activities with mistake annotations in assembly [31, 34, 42] and cooking [19, 28] domains. Based on these datasets, researchers proposed mistake detection methods [8, 12, 13, 19, 28, 33, 34]. Given videos, some researchers focus on online mistake detection using large language models (LLMs) [31] while others focus on detecting skipped steps by automatically constructing task graphs [33].

However, these studies have three limitations. First, previous work focuses on mistake detection only from videos and does not utilize procedural texts. Procedural activities are multi-modal; we read procedural texts, perceive the world, and complete steps. Therefore, to detect mistakes accurately, it is essential to input both videos and texts, rather than relying on videos alone. Second, most work analyzes whether the actions are mistakes or not. In reality, mistakes can take various forms, such as using wrong objects or skipping steps. Thus, fine-grained mistake labels are necessary to examine common mistake patterns in real-world scenarios. Finally, the target domains are limited to assembly and cooking. To gain comprehensive insights, it is essential to collect data from a wide range of domains.

To address these issues, we propose the EgoOops dataset, which consists of egocentric videos and procedural texts (Fig. 1). It covers five domains and includes 50 egocentric videos of procedural activities totaling 6.8 hours. The five domains are electrical circuits, color mixture experiments, ionic reaction experiments, toy building blocks, and cardboard crafts. Based on videos and texts, EgoOops features three types of annotations. First, we annotate the alignment between video segments (i.e., start and end timestamps) and steps in the procedural text. Second, if a segment contains a mistake, we categorize it into one of six mistake labels. Finally, for the segment with mistake labels, we provide mistake descriptions.

Based on EgoOops, we address two application tasks: video-text alignment and mistake action detection. The video-text alignment is to localize video segments and align them with corresponding steps in the procedural text. To achieve this, we utilize the self-supervised video-text alignment model StepFormer [10]. We train StepFormer on Ego4D [14] and fine-tune it on EgoOops with an additional loss, referring to this modified version as StepFormer++. The mistake action detection is to classify the segments into labels. In this paper, instead of six mistake labels, we utilize the modified three labels acquired from the original mistake labels: "mistake," "correct," and "correction." We train a multi-modal classifier to predict labels from the localized segments.

In our experiments, we prepare StepFormer and temporal action localization (TAL) models as baselines. Our approach performs better than them, achieving an F1 of 28.1 in video-text alignment and mAP of 3.7 in mistake action detection. In addition, our ablation study demonstrates the effectiveness of combining videos and texts.

2. Related work

In this section, we compare EgoOops with other datasets in terms of two perspectives: procedural activity dataset and mistake action dataset. Table 1 shows comparisons of EgoOops and other datasets.

2.1. Procedural activity dataset

Learning models from procedural activities is an important task for real-world applications, including wearable assistants [34, 39, 42, 43] and robot manipulation [1, 16, 23]. Early studies focused on alignment between video segments and steps in the procedural text and proposed datasets by collecting videos from YouTube [11, 22, 24, 38, 44, 47, 48]. Recently, egocentric videos have gained significant attention from researchers because they can record detailed workers' activities [2, 6, 14, 20, 26, 30, 37]. EPIC-KITCHEN [6] consists of 432 egocentric videos with text annotations in the cooking domain. BioVL [26] contains 16 egocentric videos with video-text alignment annotations in the wet-lab experiments. Based on these datasets, researchers develop video-text alignment models, such as ULAC [18], DWSA [35], and StepFormer [10]. EgoOops differs from these datasets because it contains not only video-text alignment but also mistake labels. This allows researchers to address both video-text alignment and mistake detection from an egocentric view.

2.2. Mistake action dataset

Mistake action detection from egocentric videos recently has gained attention from researchers, leading to the proposal of various datasets. Assembly101 [34] initially proposed a task of mistake action detection, featuring segment-level mistake annotations. ATA [13] provides video-level mistake labels across 15 classes specific to the assembly domain. HoloAssist [42] records object manipulation, where instructors verbally intervene in mistakes through a mixed-reality headset. Industreal [31] focuses on industrial-like assembly tasks and collects mistakes. Building on these datasets, previous studies have proposed various models for mistake action detection [8, 12, 33, 34, 42]. PREGO [12] addresses online mistake action detection by combining an action recognition model and LLMs to detect the mistakes. Seminara et al. [33] proposed a method to learn task graphs for detecting skipped steps by introducing a differentiable maximum likelihood loss.

Dataset	Mistakes labels			Domain	Text	Ego	#videos	Duration (hour)
	Level	Granularity	Descriptions					
YouCook2 [47]		×		Cooking	✓	×	2,000	176
CrossTask [48]		×		Diverse	×	×	4,700	375
HowTo100M [24]		×		Diverse	✓	×	1.221M	134,472
EPIC-KITCHENS [6]		×		Cooking	✓	✓	432	55
BioVL [26]		×		Biochemistry	✓	✓	16	1.6
Assembly101 [34]	Segment	Coarse	×	Assembly	×	✓	1,425	167
ATA [13]	Video	Fine	×	Assembly	×	×	1,152	24.8
HoloAssist [42]	Segment	Coarse	✓	Assembly	×	✓	2,221	166
CaptainCook4D [28]	Segment	Fine	✓	Cooking	✓	✓	384	94.5
IndustReal [31]	Segment	Coarse	✓	Assembly	✓	✓	84	5.8
EgoPER [19]	Segment	Fine	✓	Cooking	✓	✓	386	28
EgoOops (Ours)	Segment	Fine	✓	Diverse	✓	✓	50	6.8

Table 1. Comparison of EgoOops and other datasets. EgoOops is the first dataset comprising video-text alignment, mistake labels, and their descriptions in diverse domains. Level: mistake labels to a video or each temporal segment with start and end timestamps; Granularity: coarse-grained (correct/correction/mistake) or fine-grained mistake classes; Descriptions: descriptions explaining why each segment is incorrect; Text: videos accompanied by textual descriptions of actions.

These approaches focused on mistake detection only from videos and did not utilize procedural texts. However, procedural activities are multi-modal, thus it is natural to input both videos and texts. Although most datasets do not have video-text alignment, CaptainCook4D [28] and EgoPER [19] contain video-text alignment and mistake labels as with EgoOops. However, the domain of these studies is restricted to the cooking domain, and diverse domains are not verified. In addition, their approaches also rely on only videos, leaving the potential of texts unexplored. EgoOops is the first dataset that has video-text alignment and mistake labels in diverse domains. In addition, this paper demonstrates the effectiveness of combining both videos and texts for mistake action detection.

3. Dataset construction

Our EgoOops dataset consists of egocentric videos and procedural texts. The dataset construction contains the following four steps: (1) task selection, (2) preparation of procedural texts, (3) video recording, and (4) annotation of video-text alignment, mistake actions, and their descriptions.

3.1. Task selection

We aim to record procedural activities in diverse domains with various mistake actions when following procedural texts. In addition, we aim to collect diverse mistake categories (*e.g.*, unintentional action and working in the wrong way). Based on these criteria, we selected the following five tasks: electrical circuits (*EC*), color mixture experiments (*CM*), ionic reaction experiments (*IR*), toy block building (*BB*), and cardboard crafts (*CB*). These tasks satisfy the above criteria; the domains are diverse (*e.g.*, elec-

tronics, chemistry, assembly, and crafts) and contain various mistakes (*e.g.*, unintentionally cutting off the cardboard and using with wrong chemicals). In addition, the workers should follow procedural texts to accomplish the task, and the video duration ranges from a few to 30 minutes.

3.2. Preparation of procedural texts

We prepare procedural texts before recordings. For color mixture experiments and cardboard crafts, we search the web to collect procedural texts. For ion reaction experiments and electrical circuits, we use procedural texts attached to the out-of-the-box kits. For the toy block building, we write procedures from scratch because no resources exist on the web. Based on these procedural texts, we revise them to improve clarity (*e.g.* specify the tools to be used). Note that all procedural texts were originally written in Japanese, and we manually translated them to prepare the English versions. From the preliminary tests, we found that electrical circuits and toy block building were difficult to complete only with the texts, thus we gave images of the finished products to the workers.

3.3. Video recording

Participants and environments. Four Japanese graduate students (4 males) performed the tasks following procedural text. They work on every task 2 or 4 times, totaling 10 recordings for each task. The participants were equipped with a head-mounted camera Panasonic HX-A500 as shown in Fig. 2. It is a 30 fps video camera with 4K RGB resolution. To avoid the influence of background changes, a desk with objects, tools, and printed procedural texts is at the same indoor place at every recording session. Participants are recorded sitting for a close view of manipulated objects.

Figure 2. Egocentric head-mounted camera on participants.

Figure 3. Raw visual cues of copper (left), zinc (center), and magnesium (right) plates. Especially, zinc and magnesium are indistinguishable.

Figure 4. Micro QR code on objects. It's difficult to specifically identify the short thin cuboids (left) and the long thin cuboids (right) only from vision. Micro QR codes are useful for distinguishing them.

Micro QR codes. Detecting objects mentioned in procedural texts from videos is crucial for identifying incorrect actions. However, as shown in Fig. 3, some objects are visually indistinguishable. Previous works [19, 25, 27, 31] assumed manual annotations of the objects, yet it's costly and time-consuming. To address this issue, we attach Micro QR Codes [7] that encode the objects' names as shown in Fig. 4. In our experiments, we cannot utilize QR Codes, thus our future work is to leverage QR Codes for mistake action detection effectively.

Recording process. To record various mistake actions, participants intentionally perform mistake actions that we assume in advance (*e.g.*, skipping several procedures). Out of ten times to record videos, they try to follow correct procedures five times, and for the other five times, they try to contain the intentional mistakes. We call the former correct videos and the latter mistake videos. However, note that unintentional mistakes happened in some correct videos due to carelessness.

3.4. Annotations

We annotate video-text alignment, mistake labels, and their descriptions. One person annotates the whole dataset using the web annotation tool[1]. Our annotation process has three steps: (1) video-text alignment, (2) mistake labels, and (3) descriptions explaining why they are mistakes.

Video-text alignment. We first annotate video-text alignment by extracting start and end timestamps (i.e., segments) and mapping them to the corresponding steps. To reduce the ambiguity of segment annotations, we define the segment period as the time of grasping the step-related objects to the time of releasing them. Note that We also annotate the steps not written in a procedural text as undefined ones. The segments are matched with step labels. Note that any extra segments that do not correspond to any steps in the procedural texts are labeled as *undefined* (e.g., grasping the wrong object, correcting mistakes). Workers pause some steps and resume them after others (*split steps*), skip necessary steps (*missing steps*), and swap the order from the prescribed one (*out-of-order steps*), which are kinds of order mistakes.

Mistake label annotation. We define *mistakes* as deviations from the instructions in procedural texts. This definition leads to two types of mistakes: *order mistakes* and *execution mistakes*. Order mistakes occur when there is a discrepancy between the steps executed by workers and the steps outlined in the procedural texts. These include skipping, swapping, inserting extra steps, and splitting steps into temporally distant segments. Execution mistakes occur when the worker fails to follow an instruction. We classify them into the following six types: 1. working with wrong objects (*object*), 2. grasping wrong objects and releasing them without using (*mispick*), 3. correction of mistake actions (*correction*), 4. unintended actions (*accident*), 5. performing in the wrong way (*way*), 6. others (*others*).

Based on these criteria, if a segment indicates an execution mistake, we attach one of the six execution mistake labels. Note that we do not attach order mistake labels because the video-text alignment itself reveals the order mistakes. For example, if step 2 occurs before step 1, these steps are considered swapped. This categorization helps to analyze mistake patterns across diverse tasks. For example, we can analyze the frequency of mistakes in a specific task.

Descriptions. In addition, we attach descriptions to the segments with mistake labels to enable virtual assistants to explain why the actions are mistakes. To maintain the quality of the descriptions, a template for each mistake category minimizes the use of modifiers, subjects, and articles. The templates and detailed analysis of descriptions are described in the supplementary materials.

---

[1] Our annotation tool is described in the supplementary materials.

Task	Videos		Segments		
	#vid.	Avg. (min)	#seg.	#seg. / #vid.	Avg. (sec)
EC	10	3.2	98	9.8	15.4
CM	10	4.4	91	9.1	25.8
IR	10	5.4	95	9.5	29.7
BB	10	1.9	87	8.7	9.0
CB	10	26.1	167	16.7	86.7
Total	50	8.2	538	10.8	40.7

Table 2. Statistics of recorded videos.

Task	#steps in a proc. text	#words / #steps
EC	8.0	7.6
CM	8.0	17.0
IR	9.0	12.7
BB	7.0	18.6
CB	14.0	9.6

Table 3. Statistics of procedural texts.

Task	Execution mistake	Order mistake			
		Missing	Out-of-order	Split	Undefined
EC	21	2	1	12	6
CM	22	1	2	3	9
IR	19	2	4	1	6
BB	18	0	3	6	10
CB	12	7	21	29	4
Total	92	12	40	51	35

Table 4. Comparisons of execution and order mistakes.

Task	Mistake action class					
	1. object	2. mispick	3. correction	4. accident	5. way	6. others
EC	8	5	1	2	3	2
CM	4	8	0	2	5	3
IR	0	3	1	5	6	4
BB	2	5	5	1	5	0
CB	4	3	0	1	2	2
Total	18	24	7	11	21	11

Table 5. Statistics of mistake action labels.

4. Statistics

In this section, we first report video- and text-side statistics on EgoOops, then describe mistake action statistics, and finally report the annotation agreement to clarify the quality of the annotations.

4.1. Videos and procedural texts

Videos. Table 2 shows comparisons of videos between tasks. They have different trends between the tasks in terms of video duration, segments, and the number of segments. For video duration, the longest is the cardboard task at 26.1 minutes, while the shortest is the building block task at 1.9 minutes. The longest segment duration is also in the cardboard task at 87.7 seconds, while the shortest is in the building block task at 8.9 seconds.

Procedural texts. Table 3 shows comparisons of procedural texts between tasks. They have similar trends to the video statistics in terms of the number of steps in the procedural text and the number of words per step. The task with the most instructions is the cardboard task, while the task with the longest instructions is the building block task.

4.2. Mistake actions

Table 4 shows comparisons of order and execution mistakes between tasks. $40\%(= 92/230)$ segments are execution mistakes and others are order mistakes. Among the order mistakes, we observe that splitting steps is the most frequent. Table 5 shows the counts of the labels for execution mistakes. We observe some common patterns among tasks. The most two frequent labels are grasping the wrong objects and releasing them without use (label 2) and work-

ing in the wrong way or moving (label 5). In addition, we also find that each task has unique mistake patterns (*e.g.*, unintended actions frequently happen in ionic reaction experiments while they do not appear in other tasks).

4.3. Agreement

We ensure the quality of the original annotations by testing for agreement with another annotator. Because annotating all of the samples is time-consuming, we randomly choose one out of 10 videos per task. The process takes the following two steps. For video-text alignment, the annotator newly extracts segments and maps them to the corresponding steps. Then, we ask them to annotate mistake labels and descriptions. To evaluate the annotation quality, for video-text alignment, we calculate the temporal Intersection over the Union (tIoU) of segments. Note that we ignore the undefined steps for the calculation. We first calculate tIoU for each step, average by videos, and finally average them as an aggregated score. For mistake labels and descriptions, we calculate Cohen's kappa [4] and BERTScore [46], respectively.

Results. We confirm that both annotations are high quality. For the video-text alignment, we achieve 88.8, indicating a high agreement of alignment. For mistake labels and descriptions, we achieve 86.8 Cohen's kappa and 96.3 BERTScore, ensuring the high agreement of descriptions.

5. Application

Based on the constructed EgoOops dataset, we address two application tasks: video-text alignment and mistake action detection. Note that these two tasks are dependent. The outputs of the video-text alignment are used for the mistake

Figure 5. An overview of StepFormer++.

action detection. Hence, we first describe the video-text alignment and then mistake action detection.

5.1. Problem formulation

Let an untrimmed video and corresponding procedural text be $(\mathbf{V}, \mathbf{T})$. The video is written as $\mathbf{V} = (\mathbf{v}_1, \cdots, \mathbf{v}_l, \cdots, \mathbf{v}_L)$ consists of L frames. The procedural text is written as $\mathbf{T} = (\mathbf{t}_1, \cdots, \mathbf{t}_k, \cdots, \mathbf{t}_K)$. Given $(\mathbf{V}, \mathbf{T})$, the objective is to find alignment between the video and procedural text. Specifically, for k-th step $\mathbf{t}_k$, the model finds corresponding start and end timestamps $\mathbf{y}_t = (s_k, e_k)$ in the video as: $\mathbf{Y} = \{\mathbf{y}_k = (s_k, e_k) | t \in [1, K]\}$, where $\mathbf{Y}$ represents the alignment between the video and text.

Based on the alignment output, the next task is mistake action detection. Let $\mathbf{V}' = (\mathbf{v}_{s_k}, \cdots, \mathbf{v}_{e_k})$ be extracted video frames based on k-th start and end timestamps $\mathbf{y}_k$. Our objective is to classify the segment $\mathbf{V}'$ into one of the labels $\mathbf{M} = \{m_1, \ldots, m_n, \ldots, m_N\}$, where N and m_n represents the number of labels and a type of labels. We initially used the seven labels (six labels in Sec. 3.4 plus "correct" labels) to train a model. However, our preliminary experiments revealed that the model struggled to classify the segments accurately mainly due to the lack of training samples. Hence, we group mistakes except for "3. correction" into "mistake" and define the tasks as three label classification ($N = 3$): "correct," "mistake," and "correction."

5.2. Video-text alignment

For video-text alignment, we enhance StepFormer by introducing an additional loss function, resulting in Step-Former++. We will first provide an overview of the original StepFormer and then explain how we extend it.

StepFormer. StepFormer was originally proposed for learning video-text alignment from untrimmed web videos in a self-supervised manner [10]. It was trained on HowTo100M [24], which consists of video and narration pairs. StepFormer includes Transformer [41] decoders and learnable U step queries. Given video features extracted from UniVL [21], the Transformer decoder fuses the step queries and video features, producing U contextualized vec-

tors called step slots. Using the step slots and narration vectors extracted from UniVL, a sequence-to-sequence alignment is achieved via Drop-DTW [9]. Based on this alignment, the loss is calculated as InfoNCE [40] at both local (same video-narration pairs) and global (different video-narration pairs) levels. During inference, StepFormer can detect key segments without narrations. Specifically, given video features and U step slots, Drop-DTW selects meaningful slots to describe the video and drops outlier elements.

We select StepFormer because self-supervised pre-training of StepFormer can mitigate the negative impact of the small size of EgoOops. In our experiments, instead of HowTo100M, we pre-train StepFormer on Ego4D with EgoVLPv2 [29] features to fill the domain gap between web and egocentric videos. Because Ego4D is a massive-scale egocentric video dataset accompanied by transcriptions [14], we can train StepFormer as the same training procedure as the original one.

StepFormer++. We modify StepFormer with additional loss to train it on EgoOops in a fully supervised setting by leveraging video-text alignment annotations. Figure 5 shows an overview of StepFormer++. The overall process is the same as the original StepFormer. Given $(\mathbf{V}, \mathbf{T})$, the model first extract video and text features using EgoVLPv2 instead of UniVL and obtain video $\mathbf{H}_v = (\mathbf{h}_v^1, \cdots, \mathbf{h}_v^l, \cdots, \mathbf{h}_v^L)$ and step features $\mathbf{H}_t = (\mathbf{h}_t^1, \cdots, \mathbf{h}_t^k, \cdots, \mathbf{h}_t^K)$. Then, the Transformer decoder fuses $\mathbf{H}_v$ and slot queries, producing step slots $\mathbf{S} = (\mathbf{s}_1, \ldots, \mathbf{s}_u, \ldots, \mathbf{s}_U)$. Finally, the model acquires the video-text alignment via Drop-DTW by computing a similarity matrix of $\mathbf{S}$ and $\mathbf{H}_t$. The InfoNCE loss is computed at only global (different video-narration pairs) level.

Instead of the local loss, we add a new loss to the StepFormer based on video-text alignment annotations and propose StepFormer++. Based on the selected step slots $\hat{\mathbf{S}} = (\hat{\mathbf{s}}_1, \ldots, \hat{\mathbf{s}}_k, \ldots, \hat{\mathbf{s}}_K)$ after Drop-DTW[2], we train the model to align the step slots to the start and end timestamps (s_k, e_k). Specifically, this is calculated using the InfoNCE framework:

$$\mathcal{L}_{\text{supervised}}(\hat{\mathbf{s}}_k, \mathbf{y}_k, \mathbf{V}) = -\log \frac{\sum_{j \in [s_k, e_k]} f(\hat{\mathbf{s}}_k, \mathbf{v}_j)}{\sum_l f(\hat{\mathbf{s}}_k, \mathbf{v}_l)}, \quad (1)$$

where $f(\hat{\mathbf{s}}_k, \mathbf{v}_*) = \exp(\cos(\hat{\mathbf{s}}_k, \hat{\mathbf{v}}_*))/\gamma$ and γ is a scaling temperature. We add this loss to the original ones and train the model on EgoOops.

5.3. Mistake action detection

Based on the alignment results, we tackle the mistake action detection. To achieve this, we input the video and text to a multi-modal classifier to predict the label.

---

[2]Note that the number of steps is obvious when referring to the procedural text; thus we use it for Drop-DTW.

Multi-modal classifier. Given predicted video segment $\mathbf{V}'$ and k-th step $\mathbf{t}_k$, the model predicts the label. Specifically, the model first convert $\mathbf{V}'$ and $\mathbf{t}_k$ into video $\mathbf{H}'_v = (\mathbf{h}^{s_t}_v, \ldots, \mathbf{h}^{e_t}_v)$ and text features $\mathbf{h}^k_t$ using EgoVLPv2. Then, it computes the mean of $\mathbf{H}'_v$, concatenate the averaged vector with $\mathbf{h}^k_t$, and forward it into two-layer perceptrons g with ReLU function as: $\mathbf{z}_t = g(\text{concat}(\text{mean}(\mathbf{H}'_v), \mathbf{h}^k_t))$, where $\mathbf{z}_t = (z^1_t, \cdots, z^n_t, \cdots, z^N_t)$ represents the logit per label. The model applies the argmax operation on the $\mathbf{z}_t$ and outputs the prediction label.

Training. To train the model, we use the class-balanced cross-entropy loss [5] because the frequency of the "correct" label is much higher than other labels. Specifically, using $\mathbf{z}_t$, the loss is calculated as:

$$\mathcal{L}_{\text{classification}}(\mathbf{z}_t) = -\frac{1-\beta}{1-\beta^{r_{m_n}}} \log \frac{\exp(z^{m_n}_t)}{\sum_j \exp(z^{m_j}_t)}, \quad (2)$$

where r_{m_n} is the number of training samples belonging to the label m_n, and $\beta \in [0, 1)$ is a hyperparameter. We adopt teacher forcing [15, 41] as a training strategy. Specifically, we input the ground-truth segment of the t-th step to the model, rather than the predicted ones to stabilize training.

6. Experiments

6.1. Experimental settings

Splits. Our EgoOops dataset is relatively small compared to other action mistake datasets. To ensure reliable results, we perform 5-fold cross-validation. We divide the 50 videos into a 30/10/10 split for training, validation, and testing, respectively. We report the average test-set scores using the model weights that achieve the highest performance on the validation set. To construct folds, we pay attention to the two points. First, each validation and testing fold contains one correct and one mistake video for every task, totaling 10 videos. Second, each fold consist of the same workers' videos as following the group k-fold [32]. This allows us to test the models on unseen worker's activities, minimizing the bypass possibility to learn the worker-specific features to detect mistakes.

Evaluation on video-text alignment. For evaluation on video-text alignment, we follow previous work [10, 35] and utilize frame-wise precision, recall, F1, and Mean over Frames (MoF). Precision is the ratio of correctly predicted frames to the total number of frames predicted as step segments. Recall is the ratio of correctly predicted frames to the total number of step segment frames in the ground truth. The F1 score is the harmonic mean of precision and recall. MoF represents the percentage of correctly predicted frames, including those without step labels.

Evaluation on mistake action detection. For mistake action detection, we evaluate classification performance for the segments localized by the video-text alignment model. To this end, we follow temporal action localization (TAL) [3, 6, 17, 36, 45] and evaluate mean average precision (mAP) at different tIoU thresholds from 0.1 to 0.3 by 0.1. Note that mAP computes the mean of average prevision across all mistake action classes (Sec. 5.1) except for the "correct" label because we focus on mistake detection performance. Note that we compare the mistake labels only if the predicted segments correspond to the same step as the ground truth. This approach evaluates the system appropriately in terms of both step and mistake labels.

Baselines for video-text alignment. For video-text alignment, we compare StepFormer++ with the original StepFormer [10]. We evaluate two versions. First, we use StepFormer pre-trained on Ego4D and utilize it in a zero-shot manner. We call it StepFormer (ZS). Second, we fine-tune it on EgoOops in a self-supervised manner with losses proposed in [10]. We call it StepFormer (SS).

Baselines for mistake action detection. We do not adopt existing mistake detection methods as baselines because they do not fit our settings. For instance, Assembly101 [34] and CaptainCook4D [28] assume trimmed video clips as inputs, while our task assumes untrimmed videos. The method in [33] focuses on ordering mistakes in the videos and does not address execution mistakes. EgoPED [19] is the closest to our setting as it predicts both segments and mistake labels. However, their method is based on anomaly detection, predicting binary labels of "correct" and "mistakes,"; thus it cannot predict the three classes of "correct," "mistakes," and "correction."

Therefore, instead of existing mistake action detection models, we compare our method with the recent TAL model, ActionFormer [45]. This is because TAL operates under similar conditions, where the models detect both temporal segments and their action labels. In our experiments, we train ActionFormer to predict mistake labels for the detected segments, instead of action labels as in the original settings. For a fair comparison, we apply NMS to retain as many segments as the ground truths. In addition, since our metrics require step labels, we assign them to the segments in order from the start to the end of the videos. Note that the inputs for TAL are videos only.

Implementation details. We follow the official implementation of StepFormer and utilize the same hyperparameters. We set the number of step queries to be $U = 32$ and the batch size to be 6 for fine-tuning StepFormer++ on EgoOops. The video and text feature dimension of EgoVLPv2 is $d = 4,096$. We use Drop-DTW with an 80 percentile drop cost [9] for the alignment between the step slots and video features. We set $\gamma = 0.03$ in the InfoNCE loss and $\beta = 0.9999$ in the class-balanced loss. For the mistake action detection, we train the classifier in 1,200 epochs.

Methods	F1	Prec.	Rec.	MoF
StepFormer (ZS) [10]	24.1	24.6	23.7	24.9
StepFormer (SS) [10]	26.1	26.6	25.7	27.0
StepFormer++ (ours, FS)	**28.1**	**28.4**	**27.9**	**28.1**

Table 6. Results of video-text alignment. ZS: zero-shot, SS: self-supervised, FS: fully supervised.

Methods	mAP@tIoU			
	0.1	0.2	0.3	Avg.
ActionFormer [45]	1.8	0.1	0.1	0.7
StepFormer++ w/ MLP (ours)	2.5	2.6	2.5	**2.5**
GT steps (ours, oracle)				34.7

Table 7. Results of mistake action detection. As an oracle baseline, GT steps feed the classifier with ground-truth step segments to show the upper bound on complete video-text alignment. Note that TAL is a video-only task.

Inputs to classifier		Avg. mAP	Avg. mAP (oracle)
Video	Text		
✓		0.3	4.7
✓	✓	**2.5**	**34.7**

Table 8. Ablation study of inputs to the mistake action classifier of our methods.

6.2. Quantitative results

Video-text alignment. Table 6 shows the video-text alignment results. These results indicate that StepFormer++ achieves an F1 score of 28.1. This is higher than both the zero-shot and self-supervised versions of StepFormer, which have F1 scores of 24.1 and 26.1, respectively. Therefore, we confirm the effectiveness of StepFormer++.

Mistake action detection. Table 7 shows the results of mistake action detection. The results indicate that the proposed method achieves an average mAP of 2.5, which is higher than ActionFormer's 0.7. However, when compared with the oracle that uses ground truth step segments, an mAP of 2.5 is significantly lower. This suggests that improving video-text alignment performance is essential for achieving higher performance in mistake action detection.

Ablation study. Table 8 presents the ablation study of input modalities for mistake action detection to investigate the importance of procedural texts. From the results, we observe that combining both videos and texts as input achieves an mAP of 2.5. In contrast, using only videos results in an mAP of 0.3. This demonstrates the importance of incorporating both videos and procedural texts into the classifier for effective mistake action detection.

Figure 6. Qualitative results of video-text alignment.

(a) Success case. (b) Failure case.

Figure 7. Qualitative results of mistake detection. Text boxes below show the predicted steps.

6.3. Qualitative results

Figure 6 shows video-text alignment results of StepFormer (SS) and StepFormer++. While the self-supervised StepFormer fails to align video segments and text steps, StepFormer++ can correctly find the alignment, demonstrating the effectiveness of the supervised alignment loss.

Figure 7 shows examples of success and failure cases of mistake action detection. The correct video-text alignment is essential for this task. Because the incorrect alignment hurts the model's prediction, we need to improve the alignment scores continuously.

7. Conclusion

This paper proposed the EgoOops dataset, which consists of egocentric videos of workers performing five different tasks, procedural texts, and three types of annotations: video-text alignment, mistake labels, and descriptions. The mistake labels revealed both common and unique mistake patterns across the tasks. Based on EgoOops, we addressed video-text alignment and mistake action detection, proposing StepFormer++ and a multi-modal classifier built upon it. The experimental results demonstrated that our method performs better than baselines. Our abla-

tion study showed the effectiveness of the procedural texts. Future work with EgoOops includes improving the performance on both video-text alignment and mistake action detection. We currently do not use QR Codes, but they could be instrumental in achieving this by incorporating object names into the model.

References

[1] Shikhar Bahl, Russell Mendonca, Lili Chen, Unnat Jain, and Deepak Pathak. Affordances From Human Videos as a Versatile Representation for Robotics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13778–13790, 2023. 2

[2] Siddhant Bansal, Chetan Arora, and C. V. Jawahar. My View is the Best View: Procedure Learning from Egocentric Videos. In *Proceedings of the European Conference on Computer Vision*, 2022. 2

[3] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. ActivityNet: A Large-Scale Video Benchmark for Human Activity Understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 961–970, 2015. 7

[4] Jacob Cohen. A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20(1):37–46, 1960. 5

[5] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-Balanced Loss Based on Effective Number of Samples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9268–9277, 2019. 7

[6] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. The EPIC-KITCHENS Dataset: Collection, Challenges and Baselines. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(11):4125–4141, 2021. 2, 3, 7

[7] DENSO WAVE INCORPORATED. QRcode.com. https://www.qrcode.com/en/. 4

[8] Guodong Ding, Fadime Sener, Shugao Ma, and Angela Yao. Every Mistake Counts in Assembly. *arXiv preprint arXiv:2307.16453*, 2023. 2

[9] Mikita Dvornik, Isma Hadji, Konstantinos G. Derpanis, Animesh Garg, and Allan Jepson. Drop-DTW: Aligning Common Signal Between Sequences While Dropping Outliers. In *Proceedings of the Advances in Neural Information Processing Systems*, volume 34, pages 13782–13793, 2021. 6, 7

[10] Nikita Dvornik, Isma Hadji, Ran Zhang, Konstantinos G. Derpanis, Richard P. Wildes, and Allan D. Jepson. StepFormer: Self-Supervised Step Discovery and Localization in Instructional Videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18952–18961, 2023. 2, 6, 7, 8

[11] Ehsan Elhamifar and Zwe Naing. Unsupervised Procedure Learning via Joint Dynamic Summarization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6341–6350, 2019. 2

[12] Alessandro Flaborea, Guido Maria D'Amely di Melendugno, Leonardo Plini, Luca Scofano, Edoardo De Matteis, Antonino Furnari, Giovanni Maria Farinella, and Fabio Galasso. PREGO: Online mistake detection in PRocedural EGOcentric videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18483–18492, 2024. 2

[13] Reza Ghoddoosian, Isht Dwivedi, Nakul Agarwal, and Behzad Dariush. Weakly-supervised action segmentation and unseen error detection in anomalous instructional videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10128–10138, 2023. 2, 3

[14] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, Miguel Martin, Tushar Nagarajan, Ilija Radosavovic, Santhosh Kumar Ramakrishnan, Fiona Ryan, Jayant Sharma, Michael Wray, Mengmeng Xu, Eric Zhongcong Xu, Chen Zhao, Siddhant Bansal, Dhruv Batra, Vincent Cartillier, Sean Crane, Tien Do, Morrie Doulaty, Akshay Erapalli, Christoph Feichtenhofer, Adriano Fragomeni, Qichen Fu, Abrham Gebreselasie, Cristina Gonzalez, James Hillis, Xuhua Huang, Yifei Huang, Wenqi Jia, Weslie Khoo, Jachym Kolar, Satwik Kottur, Anurag Kumar, Federico Landini, Chao Li, Yanghao Li, Zhenqiang Li, Karttikeya Mangalam, Raghava Modhugu, Jonathan Munro, Tullie Murrell, Takumi Nishiyasu, Will Price, Paola Ruiz Puentes, Merey Ramazanova, Leda Sari, Kiran Somasundaram, Audrey Southerland, Yusuke Sugano, Ruijie Tao, Minh Vo, Yuchen Wang, Xindi Wu, Takuma Yagi, Ziwei Zhao, Yunyi Zhu, Pablo Arbelaez, David Crandall, Dima Damen, Giovanni Maria Farinella, Christian Fuegen, Bernard Ghanem, Vamsi Krishna Ithapu, C. V. Jawahar, Hanbyul Joo, Kris Kitani, Haizhou Li, Richard Newcombe, Aude Oliva, Hyun Soo Park, James M. Rehg, Yoichi Sato, Jianbo Shi, Mike Zheng Shou, Antonio Torralba, Lorenzo Torresani, Mingfei Yan, and Jitendra Malik. Ego4D: Around the World in 3,000 Hours of Egocentric Video. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18973–18990, 2022. 2, 6

[15] Alex Graves. Generating Sequences With Recurrent Neural Networks. *arXiv preprint arXiv:1308.0850*, 2014. 7

[16] Jiayuan Gu, Sean Kirmani, Paul Wohlhart, Yao Lu, Montserrat Gonzalez Arenas, Kanishka Rao, Wenhao Yu, Chuyuan Fu, Keerthana Gopalakrishnan, Zhuo Xu, Priya Sundaresan, Peng Xu, Hao Su, Karol Hausman, Chelsea Finn, Quan Vuong, and Ted Xiao. RT-Trajectory: Robotic Task Generalization via Hindsight Trajectory Sketches. In *Proceedings of the International Conference on Learning Representations*, 2023. 2

[17] Haroon Idrees, Amir R. Zamir, Yu-Gang Jiang, Alex Gorban, Ivan Laptev, Rahul Sukthankar, and Mubarak Shah. The THUMOS challenge on action recognition for videos "in the wild". *Computer Vision and Image Understanding*, 155:1–23, 2017. 7

[18] Anna Kukleva, Hilde Kuehne, Fadime Sener, and Jurgen Gall. Unsupervised Learning of Action Classes With Contin-

uous Temporal Embedding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12058–12066, 2019. 2

[19] Shih-Po Lee, Zijia Lu, Zekun Zhang, Minh Hoai, and Ehsan Elhamifar. Error Detection in Egocentric Procedural Task Videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18655–18666, 2024. 2, 3, 4, 7

[20] Yin Li, Miao Liu, and James M. Rehg. In the Eye of Beholder: Joint Learning of Gaze and Actions in First Person Video. In *Proceedings of the European Conference on Computer Vision*, pages 619–635, 2018. 2

[21] Huaishao Luo, Lei Ji, Botian Shi, Haoyang Huang, Nan Duan, Tianrui Li, Jason Li, Taroon Bharti, and Ming Zhou. UniVL: A Unified Video and Language Pre-Training Model for Multimodal Understanding and Generation. *arXiv preprint arXiv:2002.06353*, 2020. 6

[22] Jonathan Malmaud, Jonathan Huang, Vivek Rathod, Nicholas Johnston, Andrew Rabinovich, and Kevin Murphy. What's Cookin'? Interpreting Cooking Videos using Text, Speech and Vision. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 143–152, 2015. 2

[23] Russell Mendonca, Shikhar Bahl, and Deepak Pathak. Structured World Models from Human Videos. In *Proceedings of the Robotics: Science and Systems*. Robotics: Science and Systems Foundation, 2023. 2

[24] Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. HowTo100M: Learning a Text-Video Embedding by Watching Hundred Million Narrated Video Clips. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2630–2640, 2019. 2, 3, 6

[25] Iftekhar Naim, Young C. Song, Qiguang Liu, Liang Huang, Henry Kautz, Jiebo Luo, and Daniel Gildea. Discriminative Unsupervised Alignment of Natural Language Instructions with Corresponding Video Segments. In Rada Mihalcea, Joyce Chai, and Anoop Sarkar, editors, *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 164–174, 2015. 4

[26] Taichi Nishimura, Kojiro Sakoda, Atsushi Hashimoto, Yoshitaka Ushiku, Natsuko Tanaka, Fumihito Ono, Hirotaka Kameko, and Shinsuke Mori. Egocentric Biochemical Video-and-Language Dataset. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3129–3133, 2021. 2, 3

[27] Taichi Nishimura, Kojiro Sakoda, Atsushi Ushiku, Atsushi Hashimoto, Natsuko Okuda, Fumihito Ono, Hirotaka Kameko, and Shinsuke Mori. BioVL2: An Egocentric Biochemical Video-and-Language Dataset. *Journal of Natural Language Processing*, 29(4):1106–1137, 2022. 4

[28] Rohith Peddi, Shivvrat Arya, Bharath Challa, Likhitha Pallapothula, Akshay Vyas, Jikai Wang, Qifan Zhang, Vasundhara Komaragiri, Eric Ragan, Nicholas Ruozzi, Yu Xiang, and Vibhav Gogate. CaptainCook4D: A dataset for understanding errors in procedural activities. *arXiv preprint arXiv:2312.14556*, 2023. 2, 3, 7

[29] Shraman Pramanick, Yale Song, Sayan Nag, Kevin Qinghong Lin, Hardik Shah, Mike Zheng Shou, Rama Chellappa, and Pengchuan Zhang. EgoVLPv2: Egocentric Video-Language Pre-training with Fusion in the Backbone. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. 6

[30] Francesco Ragusa, Antonino Furnari, and Giovanni Maria Farinella. MECCANO: A multimodal egocentric dataset for humans behavior understanding in the industrial-like domain. *Computer Vision and Image Understanding*, 235:103764, 2023. 2

[31] Tim J. Schoonbeek, Tim Houben, Hans Onvlee, Peter H. N. de With, and Fons van der Sommen. IndustReal: A Dataset for Procedure Step Recognition Handling Execution Errors in Egocentric Videos in an Industrial-Like Setting. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 4365–4374, 2024. 2, 3, 4

[32] Scikit learn. 3.1. Cross-validation: Evaluating estimator performance. https://scikit-learn/stable/modules/cross_validation.html. 7

[33] Luigi Seminara, Giovanni Maria Farinella, and Antonino Furnari. Differentiable Task Graph Learning: Procedural Activity Representation and Online Mistake Detection from Egocentric Videos. *arXiv preprint arXiv:2406.01486*, 2024. 2, 7

[34] Fadime Sener, Dibyadip Chatterjee, Daniel Shelepov, Kun He, Dipika Singhania, Robert Wang, and Angela Yao. Assembly101: A Large-Scale Multi-View Video Dataset for Understanding Procedural Activities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21064–21074, 2022. 2, 3, 7

[35] Yuhan Shen, Lu Wang, and Ehsan Elhamifar. Learning to Segment Actions from Visual and Language Instructions via Differentiable Weak Sequence Alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10151–10160, 2021. 2, 7

[36] Dingfeng Shi, Yujie Zhong, Qiong Cao, Lin Ma, Jia Li, and Dacheng Tao. TriDet: Temporal Action Detection With Relative Boundary Modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18857–18866, 2023. 7

[37] Gunnar A. Sigurdsson, Abhinav Gupta, Cordelia Schmid, Ali Farhadi, and Karteek Alahari. Actor and Observer: Joint Modeling of First and Third-Person Videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7396–7404, 2018. 2

[38] Yansong Tang, Dajun Ding, Yongming Rao, Yu Zheng, Danyang Zhang, Lili Zhao, Jiwen Lu, and Jie Zhou. COIN: A Large-Scale Dataset for Comprehensive Instructional Video Analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1207–1216, 2019. 2

[39] Jesse Thomason, Michael Murray, M. Cakmak, and Luke Zettlemoyer. Vision-and-Dialog Navigation. In *Proceedings of International Conference on Robot Learning*, 2019. 2

[40] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation Learning with Contrastive Predictive Coding. *arXiv preprint arXiv:1807.03748*, 2018. 6

[41] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In *Proceedings of the Advances in Neural Information Processing Systems*, volume 30, 2017. 6, 7

[42] Xin Wang, Taein Kwon, Mahdi Rad, Bowen Pan, Ishani Chakraborty, Sean Andrist, Dan Bohus, Ashley Feniello, Bugra Tekin, Felipe Vieira Frujeri, Joshi Neel, and Marc Pollefeys. HoloAssist: An egocentric human interaction dataset for interactive AI assistants in the real world. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20270–20281, 2023. 2, 3

[43] Benita Wong, Joya Chen, You Wu, Stan Weixian Lei, Dongxing Mao, Difei Gao, and Mike Zheng Shou. AssistQ: Affordance-centric question-driven task completion for egocentric assistant. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *Proceedings of the European Conference on Computer Vision*, pages 485–501, 2022. 2

[44] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. MSR-VTT: A Large Video Description Dataset for Bridging Video and Language. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5288–5296, 2016. 2

[45] Chen-Lin Zhang, Jianxin Wu, and Yin Li. ActionFormer: Localizing Moments of Actions with Transformers. In *Proceedings of the European Conference on Computer Vision*, volume 13664, pages 492–510, 2022. 7, 8

[46] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. BERTScore: Evaluating Text Generation with BERT. In *Proceedings of the International Conference on Learning Representations*, 2020. 5

[47] Luowei Zhou, Chenliang Xu, and Jason J. Corso. Towards automatic learning of procedures from web instructional videos. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 7590–7598, 2018. 2, 3

[48] Dimitri Zhukov, Jean-Baptiste Alayrac, Ramazan Gokberk Cinbis, David Fouhey, Ivan Laptev, and Josef Sivic. Cross-Task Weakly Supervised Learning From Instructional Videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3532–3540, 2019. 2, 3

Supplementary
EgoOops: A Dataset for Mistake Action Detection
from Egocentric Videos with Procedural Texts

Yuto Haneji[1], Taichi Nishimura[2], Hirotaka Kameko[1], Keisuke Shirai[1], Tomoya Yoshida[1],
Keiya Kajimura[1], Koki Yamamoto[1], Taiyu Cui[1], Tomohiro Nishimoto[1], Shinsuke Mori[1]
[1]Kyoto University, [2]LY Corporation
haneji.yuto.s66@kyoto-u.jp, tainishi@lycorp.co.jp,
shirai.keisuke.5y@kyoto-u.ac.jp, {kameko,forest}@i.kyoto-u.ac.jp
https://y-haneji.github.io/EgoOops-project-page/

A. Annotation Tool

Figure 1 shows an annotation web tool, which consists of four panes: video player, procedural text, annotation form, and annotation list. In the video player and procedural text panes, the annotators can watch the video and read the procedural text. For video-text alignment annotations, we implemented a feature that allows them to move one second backward or forward in the video. In the annotation form pane, the annotators attach video-text alignment, mistake labels, and descriptions. For video-text alignments, annotators select the steps corresponding to the video segments. For mistake labels, annotators choose from predefined labels. For descriptions, annotators write explanations detailing why the segments are mistakes. In the annotation list pane, the annotators can view the annotation results. They can also edit or delete them by clicking on them.

B. Detailed analysis of descriptions

Description template. Table 1 shows description templates. The descriptions are written in English based on the templates, which are created based on mistake label categories. Figure 2 shows description examples annotated using the description templates.

Verb and noun distributions. Figure 3 shows the verb/nouns distributions of descriptions. We extract verbs and nouns using spaCy[1] based on en_core_web_trf. Note that for descriptions with the mistake labeled as "1. object" or "5. way", we ignore the phrase after "but" because this phrase is not a mistake but rather a correct description.

Regarding verbs, the top two verbs, "grasp" and "use," are the words in the templates. The third most frequent verb is "put," indicating mistakes related to the location and direction of objects (e.g., "put batteries in the wrong direction" in the electrical circuits). These descriptions can be confirmed in Fig. 2a. Regarding nouns, "liquid" and "plate" are the most frequent. These mistakes are common in color mixtures and ionic reactions (e.g., "use yellow liquid but should use green liquid" in the color mixture in Fig. 2b. "Block" and "switch" are the third frequent two nouns, representing mistakes in building blocks and electrical circuits, respectively.

---

[1]https://spacy.io/

tsumiki/S1760002-processed.MP4

Figure 1. Overview of annotation tool. An annotator can watch a video (upper left) and read a procedural text (upper right). In the annotation form (middle), the annotator attaches video-text alignment, mistake labels, and descriptions. The annotations are stored in the annotation list (bottom).

Mistake label	Description
1. object	Use {wrong object} but should use {correct object}
2. mispick	Grasp {wrong object}
3. correction	Correct error in step {step number}
4. accident	{verb} (+{object})
5. way	{verb} {object} {adverb} but should {verb} it/them {adverb}
6. others	Free writing. Multiple sentences are allowed to describe multple mistakes.

Table 1. Description templates.

(a) Described as "put batteries in the wrong direction."

(b) Described as "use yellow liquid but should use green liquid".

(c) Described as "drop copper plates on the table then pick them up with fingers."

Figure 2. Examples of mistake descriptions with video segments.

(a) Distribution of verbs.

(b) Distribution of nouns.

Figure 3. Distribution of verbs and nouns in mistake descriptions.