

Supplementary: Assembly101: A Large-Scale Multi-View Video Dataset for Understanding Procedural Activities

Fadime Sener[†] Dibyadip Chatterjee[‡] Daniel Shelepov[†] Kun He[†] Dipika Singhanian[‡]
Robert Wang[†] Angela Yaos[‡]

[†] Reality Labs Research @ Meta [‡] National University of Singapore
<https://assembly101.github.io/>

This supplementary further details recording settings, annotations and experiments. Section 1 provides an overview of annotator training, our recordings and custom interface. Section 2 provides the distributions of our labels, detailed train/validation/test statistics, and an extensive set of comparisons to other related datasets. Section 3 presents the architecture and implementation details of our baselines. Finally, in Section 4, we present more results and comparisons on our baselines. Please refer to the supplementary video for more information on our recordings.

1. Recording and Annotation

1.1. Recording Rig

We built a dedicated desk-based rig to capture the sequences in this dataset. Each sequence is recorded with eight RGB cameras at 1920×1080 resolution and four monochrome cameras at 640×480 resolution. Fig. 1 shows individual camera views and sample frames.

1.2. Participants

We recruited 53 adult participants (28 male, 25 female) to record approximately one-hour sessions over the course of 18 consecutive days. Participants were recruited considering the guidelines and restrictions of COVID-19, including wearing face masks. We obtained informed consent of camera wearers for the digital capture of participants, which means digitally capturing participants’ faces and bodies. All video footage and collected annotations will be made publicly available for the research community.

1.3. Annotations

1.3.1 Annotation Interface:

We developed a custom interface for annotators to temporally locate the start- and end-frames of fine-grained action segments, see Fig. 2. Each action segment is tagged with predefined verbs, tools and objects though annotators also have the flexibility for free-form entry. To promote precise

annotations, we display three static camera views to ensure that the actions are visible at least from one view without self-occlusion from the working hand. Additionally, we provide diagrams for the annotators with labelled objects of all 101 toys to ensure correct naming and terminology.

1.3.2 Annotator Training:

To ensure high-quality labels, we trained annotators over the course of four days. During this time, the annotators were introduced to our interface and the labelling task under the authors’ guidance. After training, annotators who were slow or made many mistakes were not selected to continue. Aside from training, the labelling was completed by 21 annotators over 213 hours of work.

2. Dataset Statistics & Splits

2.1. Fine-Grained Actions

From our 15 toy categories, we define 84 unique objects. Additionally, we define 24 verbs. Six of the 24 verbs are “attempted” verbs, *i.e.* the participants adjust or change their minds during assembly. For example, the “pick up chassis” action is composed of three stages of *reaching for the chassis*, *grasping it*, and *lifting it up*. Our annotators were provided with the stages of each verb. When users do not complete all stages in a segment, *e.g.* approach and/or grasp the chassis but do not lift it, we asked our annotators to place “attempt to” in front of the action. The objects and verbs combined form a total of 1456 fine-grained action labels as not every possible combination is observed. We present the distribution of our verbs and objects in Fig. 3.

To highlight the scale of our dataset, we compare the Assembly101 to other video datasets for action recognition in Table 1. Our dataset is the largest in number of segments and is richest in terms of multi-view recordings both from third-person and egocentric views.

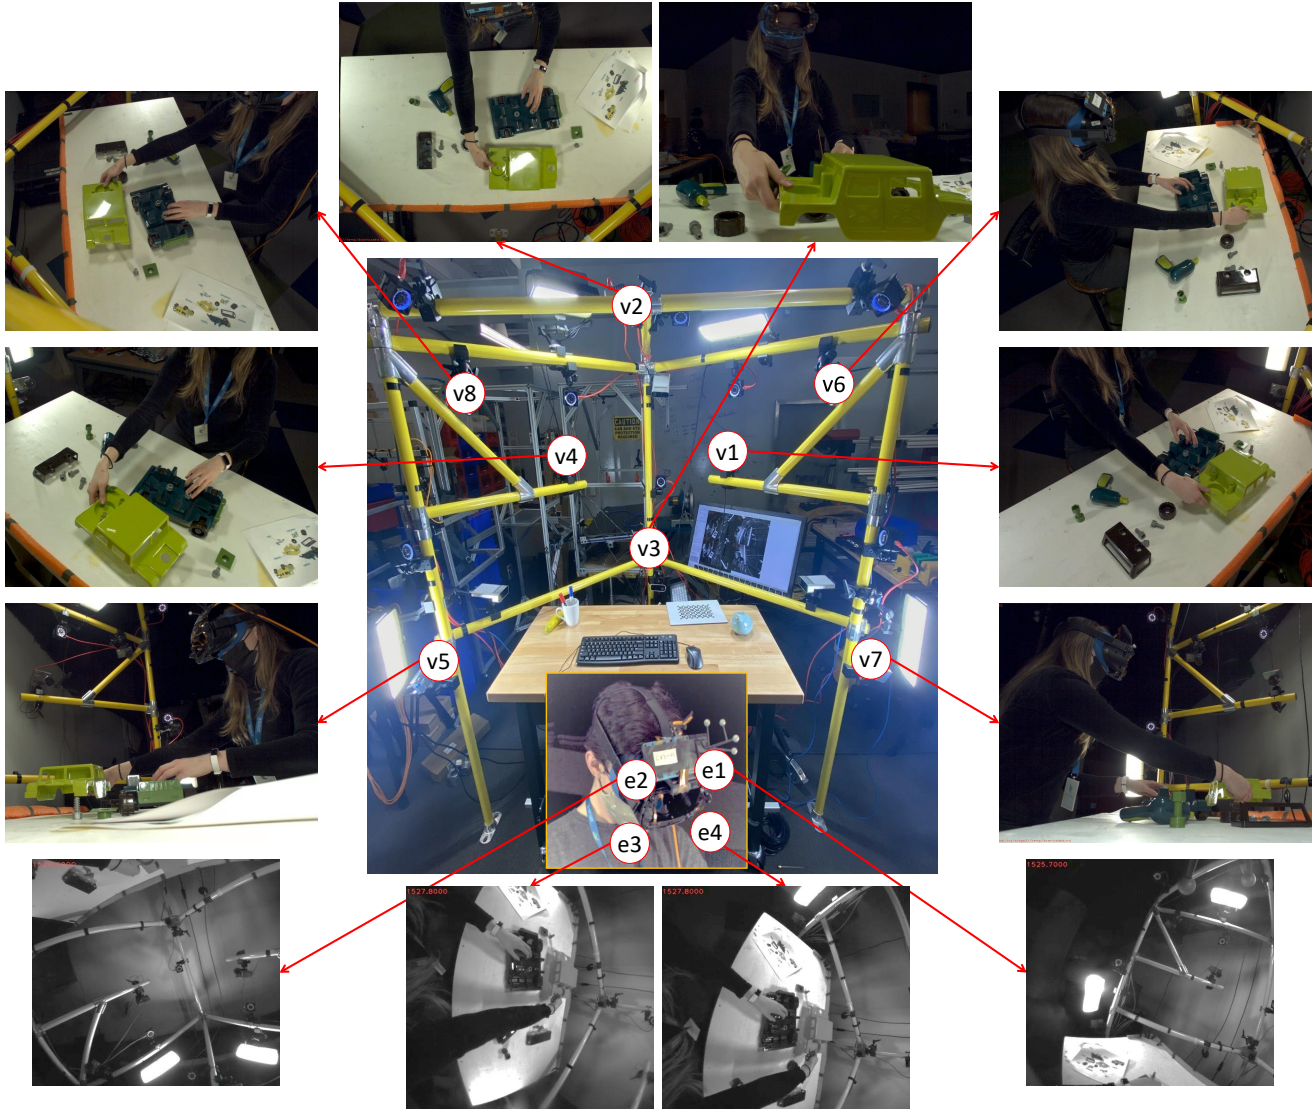


Figure 1. Our desk-based rig and sample frames from eight RGB and four monochrome cameras.

Balancing the Head: The objects, verbs, and fine-grained action labels each naturally form a long-tailed distribution [34]. When reviewing Meccano and IKEA, we observe that a handful of head-classes dominates the action distribution (60% of actions belong to 3 head classes in IKEA, and 30% in Meccano). To mitigate similar effects, we make two labelling design choices concerning the wheels and screws and tools as they are the most commonly occurring object parts. The adjustments spreads the head-tail distribution (the top 3 classes account for only 13% of action segments) and adds semantic richness to the dataset:

- **Enumerating the wheels;** *i.e.* ‘position first wheel’ vs. the generic ‘position wheel’ action. Enumeration also extends the range of temporal dependencies in a

sequence, as algorithms must keep track of how many wheels have been attached or removed.

- **Fine-grained tool and screw verbs:** Due to the nature of the assembly task, tools and screws appear very frequently. To spread head classes that result from treating tools and screws as simple objects or parts, we introduced dedicated verbs, *e.g.* “screw [object] with drill”, “position screw on [object]” and “remove screw from [object]”. Coupling these verbs with other objects conveys more information than “screw chassis” or “position screw”.

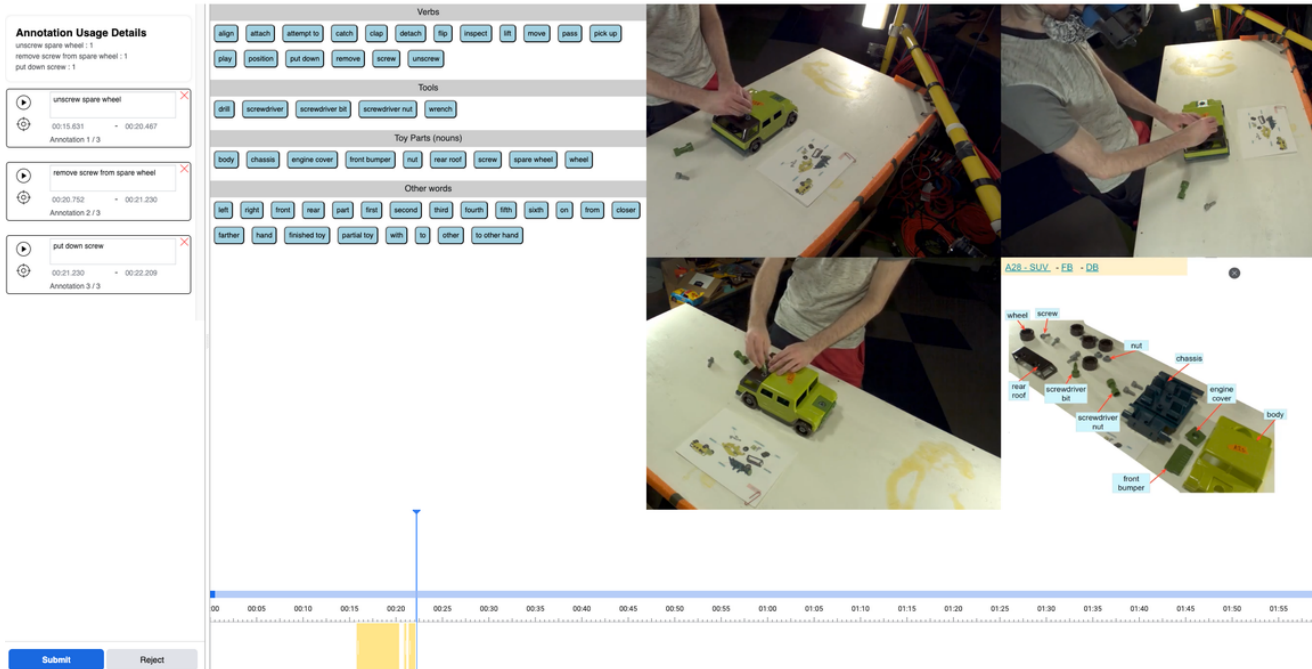


Figure 2. Our annotation tool interface. **Right panel:** input video composed of three static camera views and a diagram of the objects. **Middle panel:** pre-defined lists of verbs, tools and objects for labelling segments. The annotators also have the flexibility for free-form entry. **Bottom bar:** this annotation bar shows the temporal boundaries of the actions i.e., the start and end of each action. **Left panel:** list of temporally annotated actions.

Table 1. Comparison to other video datasets for action recognition on fine-grained actions.

Dataset	total hours	# videos	# segments	# actions	recorded	multi-view	egocentric	#pose annotation	year
MPII [28]	8.3	44	5,609	64	✓	✗	✗	✓	2012
ActivityNet [3]	648.0	27,811	23,064	200	✗	✗	✗	✗	2015
Charades [33]	81.1	9,848	67,000	157	✓	✗	✗	✗	2016
THUMOS [14]	30.0	5,613	6,310	101	✗	✗	✗	✗	2017
Charades-EGO [32]	68.8	2,751	30,516	157	✓	✓	✓	✗	2018
EPIC-100 [6]	100.0	700	89,977	4053	✓	✗	✓	✗	2020
H2O [17]	5.5	186	934	11	✓	✓	✓	✓	2021
Meccano [27]	6.9	20	8,858	61	✓	✗	✓	✗	2021
IKEAASM [2]	35.0	371	17,577	33	✓	✓	✗	✓	2021
Ego4D [12]	116.0	-	77,002	-	✓	✗	✓	✗	2021
Assembly101	517.0	4,344	991,704	1456	✓	✓	✓	✓	2021

2.2. Coarse Actions

Each coarse action is defined by the assembly or disassembly of a vehicle part. There are 193 coarse actions composed of 11 verbs and 70 objects. Each video sequence features an average of 24 coarse actions. There is an average of 10 fine-grained actions per coarse action segment. We present the distribution of coarse actions in Fig. 4. The average coarse actions in assembly sequences is 14, while it is 10 in disassembly sequences. Table 2 compares Assembly101 to other video datasets with coarse labels. Our dataset is the largest in hours and number of segments and

the only non-cooking recorded dataset.

2.3. 3D Hand Poses

Action recognition from 3D hand poses is much less explored compared to the full human body. The only existing datasets [11, 17] that focus on hand-object action recognition with 3D hand pose annotations are small scale and/or include a single hand [11]. We present our comparisons in Table 3. Compared to FPFA [11] and H2O [17], our dataset includes 82x more action segments and 200x more frames. We also compare the scale of our dataset to NTU

Table 2. Coarse action label dataset comparisons.

Dataset	hours	#videos	#segments	#actions	#recorded	#multi-view	#egocentric	#cooking	#year
GTEA [8]	0.4	28	500	71	✓	✗	✓	✓	2011
50Salads [37]	4.5	50	899	17	✓	✗	✗	✓	2013
Breakfast [16]	77	1,712	11,300	48	✓	✓	✗	✓	2014
YouTube Instructional [1]	7	150	1,260	47	✗	✗	✗	✗	2016
COIN [38]	476	11,800	46,000	778	✗	✗	✗	✗	2019
CrossTask [42]	374	4,700	34,000	107	✗	✗	✗	✓	2019
YouCookII [41]	176	2,000	15,400	-	✗	✗	✗	✓	2018
Assembly101	517	4,344	97,884	193	✓	✓	✓	✗	2021




Figure 3. We define 84 objects (**upper**) and specify 24 verbs (**bottom**), which combined with the objects form a total of 1456 fine-grained action labels. The verb distribution also shows the number of actions containing that verb on top of each bar.

Table 3. Comparisons to other datasets with 3D hand pose.

Dataset	Hours	#frames	#segments	#actions
NTU RGB+D 60 [30]	-	4M	56K	60
NTU RGB+D 120 [21]	-	8M	114K	120
FPHA [11]	1h	0.1M	1K	45
H2O [17]	5.5h	0.5M	1K	36
Assembly101	517h	111M	82K	1456

RGB+D 60 [30] and NTU RGB+D 120 [21], which are the largest full-body pose dataset. Our dataset contains 6-12x more action classes and 27-13x more frames. Additionally, NTU RGB+D 60 and NTU RGB+D 120 are composed of short trimmed clips of actions while our segments are related to each other with sequence dynamics which allows for studying the importance of temporal context for action recognition.




Figure 4. The distribution of coarse actions. Each coarse action is defined by the assembly or disassembly of a vehicle part. There are 193 coarse actions composed from 11 verbs and 70 objects.

2.4. Training, Validation & Test Splits

We use a 60/15/25 split for dividing our dataset into training, validation and test splits, with detailed statistics presented in Table 4. We present the distribution of the mis-

take action in Table 5.

For evaluation purposes, we hold out the ground truth annotations of the test split. These will be used for online challenge leaderboards to track future progress on our target

Table 4. Statistics of Assembly101 and its Train/Validation/Test splits.

Split	Hours	#videos	%	#unseen toys	#shared toys	#fine segments	#fine verbs	#fine objects	#fine actions	#coarse segments	#coarse verbs	#coarse objects	#coarse actions
Train	289	2580	60%	42	23	577,908	24	78	1251	56,916	11	65	171
Validation	80	660	15%	14	18	152,340	24	70	916	15,048	11	55	132
Test	132	1104	25%	20	19	261,456	24	79	1096	25,824	11	64	156
Overall	517	4344	100%	76	25	991,704	24	84	1456	97,884	11	70	193




Figure 5. The distribution of skill level of the participants from 1 (worst) to 5 (best). Overall, 9% of the sequences are from the participants with the worst skill level and 41% is from the best. ‘tr’, ‘v’ and ‘t’ stand for the training, validation and test splits.

Table 5. The distribution of {‘correct’, ‘mistake’, ‘correction’} segments on the coarse action segments of the assembly sequences.

	#correct	#mistake	#correction
Test	11412	2268	1992
Validation	6708	996	768
Train	25032	3768	3396
Overall	43152	7032	6156

tasks. Our dataset is designed to assess the generalizability to new toys, vehicles, actions and the participants’ skills. We thus structured our validation and test sets to examine models under varying conditions.

Seen/Unseen vehicles/toys: Of the 101 toys, only 25 toys are shared across all the three splits. We design the splits to ensure that there are also unseen vehicle categories and instances in the training to facilitate zero-shot learning. Specifically, the ‘‘SUV’’ category is not a part of the training split, while there are 20 and 14 unseen toy instances in the validation and test splits respectively.

Head vs. tail classes: The distribution of our objects and verbs can be seen in Figure 3 that has a large num-

ber of common manipulation verbs such as ‘‘pick up’’ and ‘‘put down’’ and naturally depicts a long tail distribution. The object and action distribution follows the same general trend. We define the tail classes as the set of actions classes whose instances account for 30% of the training data. This amounts to 1277 (84%) tail action classes. We used Epic-Kitchens as a reference when forming our tail classes, where 87% of the action classes are in the tail.

Skill level: Skill assessment is a critical task in many areas including sports [23], robot learning [36], surgery [39] and assembly line [25]. Which participant has the highest assembly skills? How are the participants progressing with more assembly tasks? What are the common mistakes of participants? Answering these questions involves determining how well the assembly was carried out. We, thus, also annotated the skill levels of the participant in each video from 1 (worst) to 5 (best). Overall, the distribution of skill labels in our sequences is 9%, 9%, 14%, 27% and 41% from worst to best, see Fig. 5.

3. Implementation Details

We define four action challenges: recognition, anticipation, temporal segmentation, and mistake recognition. We will release the code and evaluation scripts for all these challenges. For more robust evaluation, we will create online submission leaderboards to track the community’s progress on the test data for which the ground truth will be withheld.

3.1. Action Recognition

3.1.1 Appearance-based Action Recognition

Top-performing video based action recognition models [4, 9] are typically extensions of state-of-the-art image-based architectures [13]. Some works extend convolution and pooling to the time dimension [4, 9]; others perform channel shifting [7, 19] to capture temporal relationships while maintaining the complexity of a 2D CNN. We adopted two SOTA models as baseline for this task: TSM [19] and SlowFast [9].

Implementation Details We use two versions of the standard TSM architecture with a ResNet-50 [13] backbone - one with a single classifier head for predicting the actions

Table 6. MS-TCN++ trained on different skill levels

Trained on		F1@{10,25,50}			Edit	MoF
Skills 4,5	all	23.4	20.1	14.3	24.6	27.9
Skills 1,2,3,4,5	all	26.3	23.0	16.3	25.5	30.8

and another with two classifier heads for predicting the objects and verbs separately. Both models are trained using SGD with a momentum of 0.9, weight decay of 0.0005, and dropout of 0.5 for 50 epochs with a batch size of 64. The learning rate initialized as 0.001 is decayed by a factor of 10 at epochs 20 and 40. The best-performing model is selected via early-stopping over the validation set.

For SlowFast, we tweak the publicly available PyTorch model to have two classifier heads to predict objects and verbs separately. We use SlowFast 8x8 with a ResNet-50 backbone. The model is trained using SGD with a momentum of 0.9, weight decay of 0.0001, and dropout of 0.5 for 30 epochs with a batch size of 64. The learning rate initialized as 0.01 is decayed by a factor of 10 at epochs 20 and 25. The best-performing model is selected via early-stopping over the validation set.

3.1.2 Pose-Based Action Recognition:

State-of-the-art methods for recognizing skeleton based actions are based on deep architectures such as CNNs [20], transformers [26] and graph convolutional networks (GCN) [22, 40]. We use two SOTA GCN based methods for our experiment, 2s-AGCN [31] and MS-G3D [22].

Implementation Details We use the publicly available PyTorch [24] code for 2s-AGCN and MS-G3D. All hand pose sequences are padded to $T = 200$ frames by replaying the action segments. If there is one hand missing, we pad the second hand with 0. No data augmentation is used.

We trained 2s-AGCN [31] using Stochastic gradient descent (SGD) with Nesterov momentum (0.9) and a learning rate of 0.1 with a batch size of 32 for 30 epochs. The weight decay is set to 0.0001. For MS-G3D [22], we used SGD with a momentum of 0.9 and a learning rate of 0.05. We set the batch size to 16 and the weight decay to 0.0005. The model is trained for 50 epochs.

3.2. Action Anticipation

In our experiments, the anticipation task is defined as predicting the upcoming fine-level actions *1 second* before they start. We adopted TempAgg [29] as baseline for this task. Similar to previous works [6, 10], we report class-mean Top-5 recall as it accounts for uncertainty in future predictions.

Implementation Details We use the TempAgg with three classification heads that predicts objects, verbs and actions separately. Since TempAgg operates on frame features, we use the 2-D backbone of the TSM fine-tuned on our dataset to extract the 2048-D frame features. The *spanning past* snippet features are computed over a period of 6 seconds before the start of the action and aggregated at 3 temporal scales $K = \{5, 3, 2\}$. The *recent past* snippet features are computed over a period of $\{1.6, 1.2, 0.8, 0.4\}$ before the start of the action and aggregated over a single temporal scale $K_R = 2$. The model is trained using an Adam [15] optimizer for 15 epochs with a batch size of 32. A dropout factor of 0.3 is used. The learning rate initialised as 0.0001 is decayed by a factor of 10 after the 10th epoch.

3.3. Temporal Action Segmentation

For temporal action segmentation, we apply two competing state-of-the-art temporal convolutional networks: MS-TCN++ [18], which maintains a fixed temporal resolution in its feed-forward structure with successively larger kernel dilation, and C2F-TCN [35], a U-net-style shrink-then-expand encoder-decoder architecture. For C2F-TCN we use implicit ensembling of decoder layers and the feature augmentation strategy detailed in the paper. Performance is evaluated by mean frame-wise accuracy (MoF). Since longer actions dominate this score and it does not penalize over-segmentation errors explicitly, we also report segment-wise edit distance (Edit) and F1 scores at overlapping thresholds of 10%, 25%, and 50%, denoted as by F1@10, 25, 50.

Implementation Details Both for C2F-TCN [35] and MS-TCN++ [18] we use an Adam [15] optimizer with a batch size of 20 for a maximum of 200 epochs while using early-stopping to select the model that best fits the validation data. Loss functions used for both models are frame-wise cross entropy loss weighted with 1 and mean-square error loss [18] weighted with 0.17. For MS-TCN++, learning rate used is 0.0005 and 0 weight-decay. For C2F-TCN, learning rate of 0.001 and weight decay of 0.0001 is used. The base window for feature augmentation sampling is set to be 20 and all layers of decoder are included in ensembling.

3.4. Mistake Detection

We introduce the new problem of mistake detection in assembly videos. We adopted TempAgg [29] as the baseline for this task, which captures long-range relationships that span the order of several minutes successfully.

Implementation Details We modify the TempAgg model to capture even longer-range relationships. More precisely, the *spanning past* snippet features are computed over a period of 60 seconds around the action segment i.e. $[s-60, e+60]$, aggregated at 3 temporal scales $K = \{5, 3, 2\}$ where s

Table 7. **Action recognition** performance on fine-level actions. Presented are Top-1 accuracies.

		Overall			Head			Tail			Seen Toys			Unseen Toys		
Overall		verb	object	act.	verb	object	act.	verb	object	act.	verb	object	act.	verb	object	act.
TSM	actions	59.0	45.4	35.0	64.1	58.2	46.3	46.8	14.9	7.9	58.0	50.4	37.6	59.3	44.2	34.3
TSM	object&verb	59.3	43.9	28.9	63.9	56.2	38.8	48.4	14.8	5.3	58.0	49.7	30.9	59.6	42.5	28.4
Slowfast	object&verb	60.1	43.3	28.8	64.5	56.5	39.2	49.6	11.8	4.2	59.3	47.6	30.5	60.3	42.2	28.4

Table 8. **Action recognition & anticipation** performance on fine-level actions using TSM and TempAgg. ‘Fusion’ corresponds to average-pooling the scores from multiple views.

		Overall			Head			Tail			Seen Toys			Unseen Toys		
		verb	object	act.	verb	object	act.	verb	object	act.	verb	object	act.	verb	object	act.
Recognition	Overall	59.0	45.4	35.0	64.1	58.2	46.3	46.8	14.9	7.9	58.0	50.4	37.6	59.3	44.2	34.3
	Fusion	70.3	57.1	47.2	76.7	73.6	62.9	55.1	18.1	10.1	69.8	63.9	51.1	70.4	55.5	46.3
Anticipation	Overall	38.4	28.2	6.3	44.7	54.3	25.7	34.6	26.9	3.6	37.9	47.5	13.5	38.7	20.8	5.0
	Fusion	41.7	30.1	7.6	48.8	62.5	31.7	37.2	28.4	3.5	39.9	53.5	17.8	42.3	22.3	6.2

and e are the start and end timestamps of the action in seconds. The *recent past* snippet features are computed over a period of $\{3.0, 2.0, 1.0, 0.0\}$ around the action segment and aggregated over a single temporal scale $K_R = 5$. The training scheme remains similar to anticipation i.e. it is trained on 2048-D TSM features using an Adam [15] optimizer for 15 epochs with a batch size of 32 and a dropout of 0.3 on a single GPU. The learning rate initialised as 0.0001 is decayed by a factor of 10 after the 10th epoch.

4. Results

4.1. Action Recognition & Anticipation

In Table 7, we compare TSM trained with an action head to a TSM trained with object and verb heads. Overall, the TSM directly trained for predicting actions performs better. Although training with separate heads is expected to improve few-shot recognition [5], we observe a 2.6% decrease compared to using the action head. We also compare the performance of TSM to SlowFast.

We compare our ‘Overall’ performance to results obtained by fusing scores from multiple views on recognition and anticipation in Table 8. The fusion increases the performance of recognition significantly, while the improvement is smaller for anticipation.

4.2. Skill level

We train MS-TCN++ only on sequences with skill levels ‘4 and 5’ and then test on ‘1, 2 and 3’ sequences (see Table 6). This ablation targets at assessing whether a segmentation model trained on perfect assembly instances can generalize to sequences with mistakes. Compared to training on ‘all’ skill levels, this variant decreases MoF by 3%, indicating moderate generalizability.

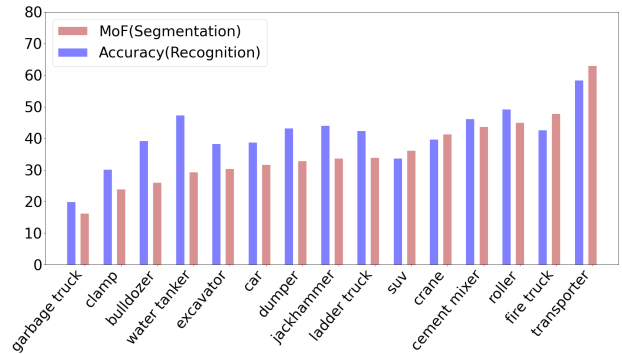


Figure 6. Action recognition accuracy and segmentation MoF over toy categories.

We did not observe a significant difference across skill levels for action recognition and anticipation tasks. We think that this is because those tasks are trained on fine-level labels while skill is more relevant for coarse actions.

4.3. Toy categories

Figure 6 shows the accuracy of action recognition and temporal action segmentation models, over toy categories individually. The toy with the highest score is ‘transporter’. Although we have only 4 toys in ‘transporter’ category, we have 22 participants recording these toys. We think its high performance could be due to the large number of recordings. ‘SUV’ category only appears in the test set. However, the performance on this toy is still competitive at 33.5%, indicating the success of generalization from other categories such as cars.




Figure 7. Action recognition object and verb recall.

4.4. Class-based evaluations

Fine-grained Actions We present the recall of the object and verbs for action recognition in Fig. 7. The verbs with the highest recall are ‘clap’, ‘pick up’ and ‘put down’, while the tail verbs of ‘attempt to’ have the lowest recall. We present the top 24 object classes as well. It can be seen that enumerated wheels are among the top classes.

Coarse Actions Based on temporal action segmentation results, we further investigated the performance of verbs and objects. Out of 11 coarse verbs, the verbs with the highest recall are ‘demonstrate’, ‘attach’ and ‘detach’, and the ones with the lowest recall are ‘position’, ‘remove’ and ‘attempt to screw’, which are the tail verbs. The objects with the highest recall are ‘chassis’ and ‘interior’ which are the most common objects across toys.

References

[1] Jean-Baptiste Alayrac, Piotr Bojanowski, Nishant Agrawal, Josef Sivic, Ivan Laptev, and Simon Lacoste-Julien. Unsupervised learning from narrated instruction videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4575–4583, 2016. 4

[2] Yizhak Ben-Shabat, Xin Yu, Fatemeh Saleh, Dylan Campbell, Cristian Rodriguez-Opazo, Hongdong Li, and Stephen Gould. The ikea asm dataset: Understanding people assembling furniture through actions, objects and pose. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 847–859, 2021. 3

[3] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceed-*

ings of the IEEE conference on computer vision and pattern recognition, pages 961–970, 2015. 3

[4] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6299–6308, 2017. 6

[5] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Scaling egocentric vision: The epic-kitchens dataset. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 8

[6] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Antonino Furnari, Jian Ma, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Rescaling egocentric vision. *CoRR*, abs/2006.13256, 2020. 3, 7

[7] Linxi Fan*, Shyamal Buch*, Guanzhi Wang, Ryan Cao, Yuke Zhu, Juan Carlos Niebles, and Li Fei-Fei. Rubiknet: Learnable 3d-shift for efficient video action recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 6

[8] Alireza Fathi, Xiaofeng Ren, and James M Rehg. Learning to recognize objects in egocentric activities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011. 4

[9] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6202–6211, 2019. 6

[10] Antonino Furnari, Sebastiano Battiato, and Giovanni Maria Farinella. Leveraging uncertainty to rethink loss functions and evaluation measures for egocentric action anticipation. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018. 7

[11] Guillermo Garcia-Hernando, Shanxin Yuan, Seungryul Baek, and Tae-Kyun Kim. First-person hand action benchmark with rgb-d videos and 3d hand pose annotations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 409–419, 2018. 3, 4

[12] 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, ayant 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, Christian Fuegen, Abrahm Gebreselasie, Cristina Gonz alez, James Hillis, Xuhua Huang, Yifei Huang, Wenqi Jia, Weslie Khoo, Jachym Kolar, Satwik Kottur, Anurag Kumar, Federico Landini, Chao Li, Zhenqiang Li, Karttikeya Mangalam, Raghava Modhugu, Jonathan Munro, Tullie Murrell, Takumi Nishiyasu, Will Price, Paola Ruiz Puentes, Mery Ramazanova, Leda Sari, Kiran Somasundaram, Audrey Southerland, Yusuke Sugano, Ruijie Tao, Minh Vo,

- Yuchen Wang, Xindi Wu, Takuma Yagi, Yunyi Zhu, ablo Arbel aez, David Crandall⁶, Dima Damen, Giovanni Maria Farinella, 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. Around the World in 3,000 Hours of Egocentric Video. *CoRR*, abs/2110.07058, 2021. 3
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [14] 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. 3
- [15] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 7, 8
- [16] Hilde Kuehne, Ali Arslan, and Thomas Serre. The language of actions: Recovering the syntax and semantics of goal-directed human activities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 4
- [17] Taein Kwon, Bugra Tekin, Jan Stuhmer, Federica Bogo, and Marc Pollefeys. H2o: Two hands manipulating objects for first person interaction recognition. *arXiv preprint arXiv:2104.11181*, 2021. 3, 4
- [18] Shi-Jie Li, Yazan AbuFarha, Yun Liu, Ming-Ming Cheng, and Juergen Gall. Ms-tcn++: Multi-stage temporal convolutional network for action segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2020. 7
- [19] Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7083–7093, 2019. 6
- [20] Jian Liu, Naveed Akhtar, and Ajmal Mian. Skepxels: Spatio-temporal image representation of human skeleton joints for action recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2019. 7
- [21] Jun Liu, Amir Shahroudy, Mauricio Perez, Gang Wang, Ling-Yu Duan, and Alex C Kot. Ntu rgb+ d 120: A large-scale benchmark for 3d human activity understanding. *IEEE transactions on pattern analysis and machine intelligence*, 42(10):2684–2701, 2019. 4
- [22] Ziyu Liu, Hongwen Zhang, Zhenghao Chen, Zhiyong Wang, and Wanli Ouyang. Disentangling and unifying graph convolutions for skeleton-based action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 143–152, 2020. 7
- [23] Paritosh Parmar and Brendan Tran Morris. Learning to score olympic events. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 20–28, 2017. 6
- [24] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32:8026–8037, 2019. 7
- [25] Mikkel Rath Pedersen, Lazaros Nalpantidis, Rasmus Skovgaard Andersen, Casper Schou, Simon Bøgh, Volker Krüger, and Ole Madsen. Robot skills for manufacturing: From concept to industrial deployment. *Robotics and Computer-Integrated Manufacturing*, 37:282–291, 2016. 6
- [26] Chiara Plizzari, Marco Cannici, and Matteo Matteucci. Spatial temporal transformer network for skeleton-based action recognition. In *International Conference on Pattern Recognition*, pages 694–701. Springer, 2021. 7
- [27] Francesco Ragusa, Antonino Furnari, Salvatore Livatino, and Giovanni Maria Farinella. The meccano dataset: Understanding human-object interactions from egocentric videos in an industrial-like domain. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1569–1578, 2021. 3
- [28] Marcus Rohrbach, Sikandar Amin, Mykhaylo Andriluka, and Bernt Schiele. A database for fine grained activity detection of cooking activities. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2012. 3
- [29] Fadime Sener, Dipika Singhania, and Angela Yao. Temporal aggregate representations for long-range video understanding. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 154–171. Springer, 2020. 7
- [30] Amir Shahroudy, Jun Liu, Tian-Tsong Ng, and Gang Wang. Ntu rgb+ d: A large scale dataset for 3d human activity analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1010–1019, 2016. 4
- [31] Lei Shi, Yifan Zhang, Jian Cheng, and Hanqing Lu. Two-stream adaptive graph convolutional networks for skeleton-based action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12026–12035, 2019. 7
- [32] 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/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7396–7404, 2018. 3
- [33] Gunnar A. Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016. 3
- [34] Eero P Simoncelli and Bruno A Olshausen. Natural image statistics and neural representation. *Annual review of neuroscience*, 24(1):1193–1216, 2001. 2
- [35] Dipika Singhania, Rahul Rahaman, and Angela Yao. Coarse to fine multi-resolution temporal convolutional network. *arXiv preprint arXiv:2105.10859*, 2021. 7

- [36] Bradly C Stadie, Pieter Abbeel, and Ilya Sutskever. Third-person imitation learning. *arXiv preprint arXiv:1703.01703*, 2017. 6
- [37] Stein and McKenna. Combining embedded accelerometers with computer vision for recognizing food preparation activities. In *UbiComp*, 2013. 4
- [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 (CVPR)*, 2019. 4
- [39] S Swaroop Vedula, Anand Malpani, Narges Ahmidi, Sanjeev Khudanpur, Gregory Hager, and Chi Chiung Grace Chen. Task-level vs. segment-level quantitative metrics for surgical skill assessment. *Journal of surgical education*, 73(3):482–489, 2016. 6
- [40] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Thirty-second AAAI conference on artificial intelligence*, 2018. 7
- [41] Luowei Zhou, Chenliang Xu, and Jason J Corso. Towards automatic learning of procedures from web instructional videos. In *AAAI*, 2018. 4
- [42] 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 (CVPR)*, 2019. 4