

Trespassing the Boundaries: Labeling Temporal Bounds for Object Interactions in Egocentric Video

Davide Moltisanti Michael Wray
University of Bristol, United Kingdom

Walterio Mayol-Cuevas Dima Damen
<FirstName>.<LastName>@bristol.ac.uk

Abstract

Manual annotations of temporal bounds for object interactions (i.e. start and end times) are typical training input to recognition, localization and detection algorithms. For three publicly available egocentric datasets, we uncover inconsistencies in ground truth temporal bounds within and across annotators and datasets. We systematically assess the robustness of state-of-the-art approaches to changes in labeled temporal bounds, for object interaction recognition. As boundaries are trespassed, a drop of up to 10% is observed for both Improved Dense Trajectories and Two-Stream Convolutional Neural Network.

We demonstrate that such disagreement stems from a limited understanding of the distinct phases of an action, and propose annotating based on the Rubicon Boundaries, inspired by a similarly named cognitive model, for consistent temporal bounds of object interactions. Evaluated on a public dataset, we report a 4% increase in overall accuracy, and an increase in accuracy for 55% of classes when Rubicon Boundaries are used for temporal annotations.

1. Introduction

Egocentric videos, also referred to as first-person videos, have been frequently advocated to provide a unique perspective into object interactions [12, 5, 19]. These capture the viewpoint of an object close to that perceived by the user during interactions. Consider, for example, ‘turning a door handle’. Similar appearance and motion information will be captured from an egocentric perspective as multiple people turn a variety of door handles.

Several datasets have been availed to the research community focusing on object interactions from head-mounted [2, 7, 6, 1, 15] and chest-mounted [21] cameras. These incorporate ground truth labels that mark the *start* and the *end* of each object interaction, such as ‘open fridge’, ‘cut tomato’ and ‘push door’. These temporal bounds are the base for automating action detection, localization and recognition. They are thus highly influential in the ability

of an algorithm to distinguish one interaction from another.

As temporal bounds vary, the segments may contain different portions of the untrimmed video from which the action is extracted. Humans can still recognize an action even when the video snippet varies or contains only part of the action. Machines are not yet as robust, given that current algorithms strongly rely on the data and the labels we feed to them. Should these bounds be incorrectly or inconsistently annotated, the ability to learn as well as assess models for action recognition would be adversely affected.

In this paper, we uncover inconsistencies in defining temporal bounds for object interactions within and across three egocentric datasets. We show that temporal bounds are often ill-defined, with limited insight into how they have been annotated. We systematically show that perturbations of temporal bounds influence the accuracy of action recognition, for both hand-crafted features and fine-tuned classifiers, even when the tested video segment significantly overlaps with the ground truth segment.

While this paper focuses on unearthing inconsistencies in temporal bounds, and assessing their effect on object interaction recognition, we take a step further into proposing an approach for consistently labeling temporal bounds inspired by studies in the human mindset.

Main Contributions More specifically, we:

- Inspect the consistency of temporal bounds for object interactions *across and within* three datasets for egocentric object interactions. We demonstrate that current approaches are highly subjective, with visible variability in temporal bounds when annotating instances of the same action;
- Evaluate the robustness of two state-of-the-art action recognition approaches, namely Improved Dense Trajectories [32] and Convolutional Two-Stream Network Fusion [8], to changes in temporal bounds. We demonstrate that the recognition rate drops by 2-10% when temporal bounds are modified albeit within an Intersection-over-Union of more than 0.5;
- Propose, inspired by studies in Psychology, the Rubi-

con Boundaries to assist in consistent temporal boundary annotations for object interactions;

- Re-annotate one dataset using the Rubicon Boundaries, and show more than 4% increase in recognition accuracy, with improved per-class accuracies for most classes in the dataset.

We next review related works in Section 2, before embarking on inspecting labeling consistencies in Section 3, evaluating recognition robustness in Section 4 and proposing and evaluating the Rubicon Boundaries in Section 5. The paper concludes with an insight into future directions.

2. Related Work

In this section, we review all papers that, up to our knowledge, ventured into the consistency and robustness of temporal bounds for action recognition.

Temporal Bounds in Non-Egocentric Datasets The leading work of Satkin and Hebert [24] first pointed out that determining the temporal extent of an action is often subjective, and that action recognition results vary depending on the bounds used for training. They proposed to find the most discriminative portion of each segment for the task of action recognition. Given a loosely trimmed training segment, they exhaustively search for the cropping that leads to the highest classification accuracy, using hand-crafted features such as HOG, HOF [13] and Trajectons [18]. Optimizing bounds to maximize discrimination between class labels has also been attempted by Duchenne *et al.* [3], where they refined loosely labeled temporal bounds of actions, estimated from film scripts, to increase accuracy across action classes. Similarly, two works evaluated the optimal segment length for action recognition [25, 36]. From the *start* of the segment, 1-7 frames were deemed sufficient in [25], with rapidly diminishing returns as more frames were added. More recently, [36] showed that 15-20 frames were enough to recognize human actions from 3D skeleton joints.

Interestingly, assessing the effect of temporal bounds is still an active research topic within novel deep architectures. Recently, Peng *et al.* [20] assessed how frame-level classifications using multi-region two-stream CNN are pooled to achieve video-level recognition results. The authors reported that stacking more than 5 frames worsened the action detection and recognition results for the tested datasets, though only compared to a 10-frame stack.

The problem of finding optimal temporal bounds is much akin to that of action localization in untrimmed videos [33, 14, 11]. Typical approaches attempt to find similar temporal bounds to those used in training, making them equally dependent on manual labels and thus sensitive to inconsistencies in the ground truth labels.

An interesting approach that addressed reliance on training temporal bounds for action recognition and localization

is that of Gaidon *et al.* [9]. They noted that action recognition methods rely on temporal bounds in test videos to be strictly containing an action, and in *exactly* the same fashion as the training segments. They thus redefined an action as a sequence of key atomic frames, referred to as actoms. The authors learned the optimal sequence of actoms per action class with promising results. More recently, Wang *et al.* [34] represented actions as a transformation from a precondition state to an effect state. The authors attempted to learn such transformations as well as locate the end of the precondition and the start of the effect. However, both approaches [9, 34] rely on manual annotations of actoms [9] or action segments [34], which are potentially as subjective as the temporal bounds of the actions themselves.

Temporal Bounds in Egocentric Datasets Compared to third person action recognition (e.g. 101 action classes in [28] and 157 action classes in [26]), egocentric datasets have a smaller number of classes (5-44 classes [2, 7, 6, 1, 15, 21, 37]), with considerable ambiguities (e.g. ‘turn on’ vs ‘turn off’ tap). Comparative recognition results have been reported on these datasets in [29, 31, 27, 16, 22, 17].

Previously, three works noted the challenge and difficulty in defining temporal bounds for egocentric videos [29, 1, 37]. In [29], Spriggs *et al.* discussed the level of granularity in action labels (e.g. ‘break egg’ vs ‘beat egg in a bowl’) for the CMU dataset [2]. They also noted the presence of temporally overlapping object interactions (e.g. ‘pour’ while ‘stirring’). In [35], multiple annotators were asked to provide temporal bounds for the same object interaction. The authors showed variability in annotations, yet did not detail what instructions were given to annotators when labeling these temporal bounds. In [37], the human ability to order pairwise egocentric segments was evaluated as the snippet length varied. The work showed that human perception improves as the size of the segment increases to 60 frames, then levels off.

To summarize, temporal bounds for object interactions in egocentric video have been overlooked and no previous work attempted to analyze the influence of consistency of temporal bounds or the robustness of representations to variability in these bounds. This paper particularly attempts to answer both questions; *how consistent are current temporal bound labels in egocentric datasets?* and *how sensitive are action recognition results to changes in these temporal bounds?* We next delve into answering these questions.

3. Temporal Bounds: Inspecting Inconsistency

Current egocentric datasets are annotated for a number of action classes, described using a verb-noun label. Each class instance is annotated with its label as well as the temporal bounds (i.e. *start* and *end* times) that delimit the frames used to learn the action model. Little information is typically provided on how these manually labeled tempo-




Figure 1: Annotations for action ‘pour sugar/oil’ from BEOID, GTEA Gaze+ and CMU. Aligned key frames are shown along with ground truth annotations (red). The yellow rectangle encloses the motion strictly involved in ‘pour’.

ral bounds are acquired. In Section 3.1, we compare labels across and within egocentric datasets. We then discuss in Section 3.2 how variability is further increased when multiple annotators for the same action are employed.

3.1. Labeling in Current Egocentric Datasets

We study ground truth annotations for three public datasets, namely BEOID [1], GTEA Gaze+ [6] and CMU [2]. Observably, many annotations base the *start* and *end* of an action as respectively the first and last frames when the hands are visible in the field of view. Other annotations tend to segment an action more strictly, including only the most relevant physical object interaction within the bounds. Figure 1 illustrates an example of three different temporal bounds for the ‘pour’ action across the three datasets. Frames marked in red are those that have been labeled in the different datasets as containing the ‘pour’ action. The annotated temporal bounds in this example vary remarkably; (i) BEOID’s are the tightest; (ii) The start of GTEA Gaze+’s segment is belated: in fact, the first frame in the annotated segment shows some oil already in the pan; (iii) CMU’s segment includes picking the oil container and putting it down before and after pouring. These conclusions extend to other actions in the three datasets.

We observe that annotations are also inconsistent within the same dataset. Figure 2 shows three intra-dataset annotations. (i) For the action ‘open door’ in BEOID, one segment includes the hand reaching the door, while the other starts with the hand already holding the door’s handle; (ii) For the action ‘cut pepper’ in GTEA Gaze+, in one segment the user already holds the knife and cuts a single slice of the vegetable. The second segment includes the action of picking up the knife, and shows the subject slicing the whole pepper through several cuts. Note that the length difference between the two segments is considerable - the segments are respectively 3 and 80 seconds long; (iii) For the action ‘crack egg’ in CMU, only the first segment shows the user tapping the egg against the bowl.

While the figure shows three examples, such inconsis-




Figure 2: Inconsistency of temporal bounds within datasets. Two segments from each action are shown with considerable differences in start and end times.

tencies have been discovered throughout the three datasets. However, we generally observe that GTEA Gaze+ shows more inconsistencies, which could be due to the dataset size, as it is the largest among the evaluated datasets.

3.2. Multi-Annotator Labeling

As noted above, defining when an object interaction begins and finishes is highly subjective. There is usually little agreement when different annotators segment the same object interaction. To assess this variability, we collected 5 annotations for several object interactions from an untrimmed video of the BEOID dataset. First, annotators were only informed of the class name and asked to label the start and the end of the action. We refer to these annotations as *conventional*. We then asked a different set of annotators to annotate the same object interactions following our proposed Rubicon Boundaries (RB) approach which we will present in Section 5. Figure 3 shows per-class box plots for the Intersection-over-Union (IoU) measure for all pairs of annotations. RB annotations demonstrate gained consistency for all classes. For conventional annotations, we report an average IoU = 0.63 and a standard deviation of 0.22, whereas for RB annotations we report increased average IoU = 0.83 with a lower standard deviation of 0.11.

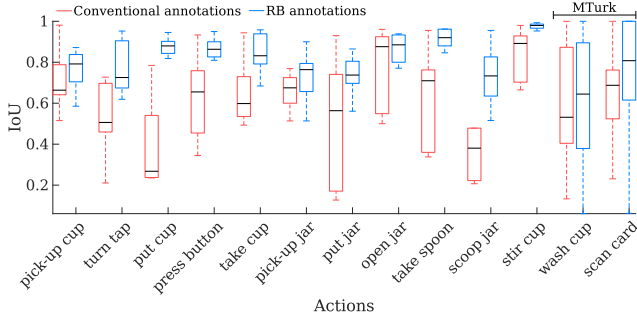


Figure 3: IoU comparison between conventional (red) and RB (blue) annotations for several object interactions.

To assess how consistency changes as more annotations are collected, we employ annotators via the Amazon Mechanical Turk (MTurk) for two object interactions from BEOID, namely the actions of ‘scan card’ and ‘wash cup’, for which we gathered 45 conventional and RB labels. Box plots for MTurk labels are included in Figure 3, showing marginal improvements with RB annotations as well. We will revisit RB annotations in detail in Section 5.

In the next Section, we assess the robustness of current action recognition approaches to variations in temporal boundaries.

4. Temporal Bounds: Assessing Robustness

To assess the effect of temporal bounds on action recognition, we systematically vary the *start* and *end* times of annotated segments for the three datasets, and report comprehensive results on the effect of such alterations.

Results are evaluated using 5-fold cross validation. For training, only ground truth segments are considered. We then classify *both* the original ground truth and the generated segments. We provide results using Improved Dense Trajectories [32] encoded with Fisher Vector [23] (IDT FV)¹ and Convolutional Two-Stream Network Fusion for Video Action Recognition (2SCNN) [8]. The encoded IDT FV features are classified with a linear SVM. Experiments on 2SCNN are carried out using the provided code and the proposed VGG-16 architecture pre-trained on ImageNet and tuned on UCF101 [28]. We fine-tune the spatial, temporal and fusion networks on each fold’s training set.

Theoretically, the two action recognition approaches are likely to respond differently to variations in start and end times. Specifically, 2SCNN averages the classification responses of the fusion network obtained on n frames randomly extracted from a test video v of length $|v|$. In our experiments, $n = \min(20, |v|)$. Such strategy should ascribe some degree of resilience to 2SCNN. IDT densely samples feature points in the first frame of the video, whereas in the following frames only new feature points are sampled to re-

¹IDT features have been extracted using GNU Parallel [30].

Dataset	N. of gt segments	N. of gen segments	Classes
BEOID [1]	742	16691	34
GTEA Gaze+ [6]	1141	22221	42
CMU [2]	450	26160	31

Table 1: Number of ground truth/generated segments and number of classes for BEOID, GTEA Gaze+ and CMU.

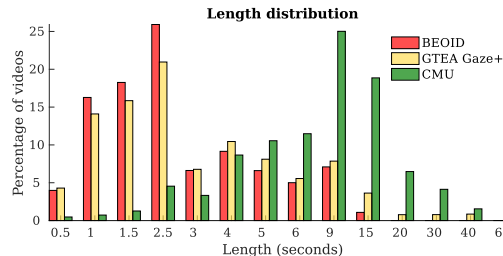


Figure 4: Video’s length distribution across datasets.

place the missing ones. This entails that IDT FV should be more sensitive to start (specifically) and end time variations, at least for shorter videos. This fundamental difference makes both approaches interesting to assess for robustness.

4.1. Generating Segments

Let v_{gt} be a ground truth action segment obtained by cropping an untrimmed video from time s_{gt} to time e_{gt} , which denote the annotated ground truth start and end times. We vary both s_{gt} and e_{gt} in order to generate new action segments with different temporal bounds. More precisely, let $s_{gen}^0 = s_{gt} - \Delta$ and let $s_{gen}^n = s_{gt} + \Delta$. The set containing candidate start times is defined as:

$$\mathcal{S} = \{s_{gen}^0, s_{gen}^0 + \delta, s_{gen}^0 + 2\delta, \dots, s_{gen}^0 + (n-1)\delta, s_{gen}^n\}$$

Analogously, let $e_{gen}^0 = e_{gt} - \Delta$ and let $e_{gen}^n = e_{gt} + \Delta$, the set containing candidate end times is defined as:

$$\mathcal{E} = \{e_{gen}^0, e_{gen}^0 + \delta, e_{gen}^0 + 2\delta, \dots, e_{gen}^0 + (n-1)\delta, e_{gen}^n\}$$

To accumulate the set of generated action segments, we take all possible combinations of $s_{gen} \in \mathcal{S}$ and $e_{gen} \in \mathcal{E}$ and keep only those such that the Intersection-over-Union between $[s_{gt}, e_{gt}]$ and $[s_{gen}, e_{gen}] \geq 0.5$. In all our experiments, we set $\Delta = 2$ and $\delta = 0.5$ seconds.

4.2. Comparative Evaluation

Table 1 reports the number of ground truth and generated segments for BEOID, GTEA Gaze+ and CMU. Figure 4 illustrates the segments’ length distribution for the three datasets, showing considerable differences: BEOID and GTEA Gaze+ contain mostly short segments (1-2.5 seconds), although the latter includes also videos up to 40 seconds long. CMU has longer segments, with the majority ranging from 5 to 15 seconds.

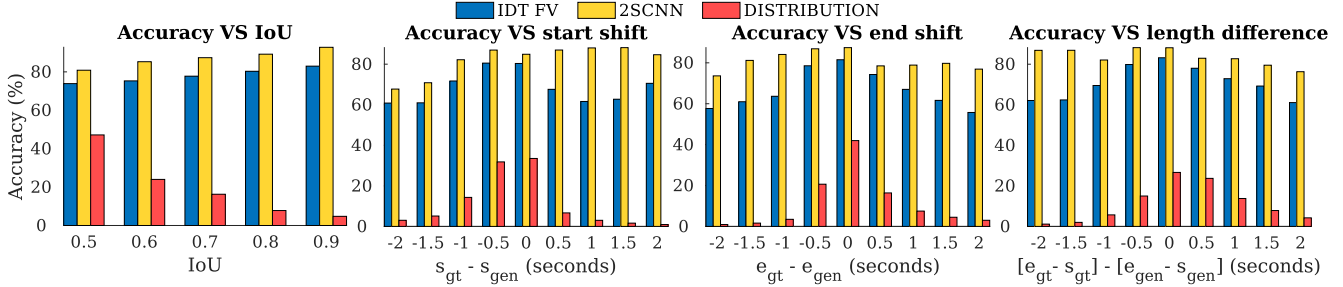


Figure 5: BEOID: classification accuracy vs IoU, start/end shifts and length difference between gt and generated segments.

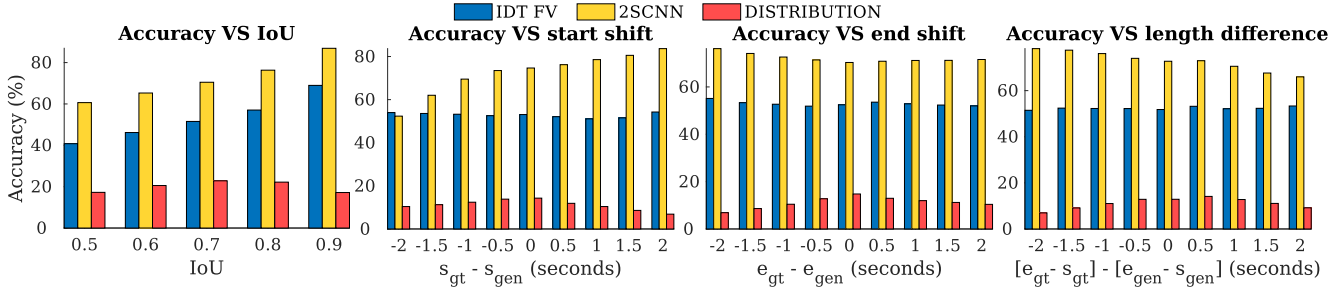


Figure 6: CMU: classification accuracy vs IoU, start/end shifts and length difference.

BEOID [1] is the evaluated dataset with the most consistent and the tightest temporal bounds. When testing the ground truth segments using both IDT FV and 2SCNN, we observe high accuracy for ground truth segments - respectively 85.3% and 93.5% - as shown in Table 2. When classifying the generated segments, we observe a drop in accuracy of 9.9% and 9.7% respectively.

Figure 5 shows detailed results where accuracy is reported vs IoU, start/end shifts and length difference between ground truth and generated segments. We particularly show the results for shifts in the start and the end times independently. A *negative start shift* implies that a generated segment begins before the corresponding ground truth segment, and a *negative end shift* implies that a generated segment finishes before the corresponding ground truth segment. These terms are used consistently throughout this section. Results show that: (i) as IoU decreases the accuracy drops consistently for IDT FV and 2SCNN - which questions both approaches' robustness to temporal bounds alterations; (ii) IDT FV exhibits lower accuracy with both negative and positive start/end shifts; (iii) IDT FV similarly exhibits lower accuracy with negative and positive length differences. This is justified as BEOID segments are tight; by expanding a segment we include new potentially noisy or irrelevant frames that confuse the classifiers; (iv) 2SCNN is more robust to length difference which is understandable as it randomly samples a maximum of 20 frames regardless of the length. While these are somehow expected, we also note that (v) 2SCNN is robust to *positive start shifts*.

Dataset	IDT FV _{gt}	IDT FV _{gen}	2SCNN _{gt}	2SCNN _{gen}
BEOID	85.3	75.4	93.5	83.8
CMU	54.9	52.8	76.0	71.7
GTEA Gaze+	45.4	43.3	61.2	59.6

Table 2: Classification accuracy for ground truth and generated segments for BEOID, CMU and GTEA Gaze+.

CMU [2] is the dataset with longer ground truth segments. Table 2 compares results obtained for CMU's ground truth and generated segments. For this dataset, IDT FV accuracy drops by 2.1% for the generated segments, whereas 2SCNN drops by 4.3%. In Figure 6, CMU consistently shows low robustness for both IDT FV and 2SCNN. As IoU changes from > 0.9 to > 0.5 , we observe a drop of more than 20% in accuracy for both. However, due to the long average length of segments in CMU, the effect of shifts in start end times as well as length differences is not visible for IDT FV. Interestingly for 2SCNN, the accuracy slightly improves with *positive start shift*, *negative end shift* and *negative length difference*. This suggests that CMU's ground truth bounds are somewhat loose and that tighter segments are likely to contain more discriminative frames.

GTEA Gaze+ [6] is the dataset with the most inconsistent bounds, based on our observations. Table 2 shows that accuracy for IDT FV drops by 2.1%, while overall accuracy for 2SCNN drops marginally (1.6%). This should not be mistaken for robustness, and that is evident when studying the results in Figure 7. For all variations (i.e. start/end shifts

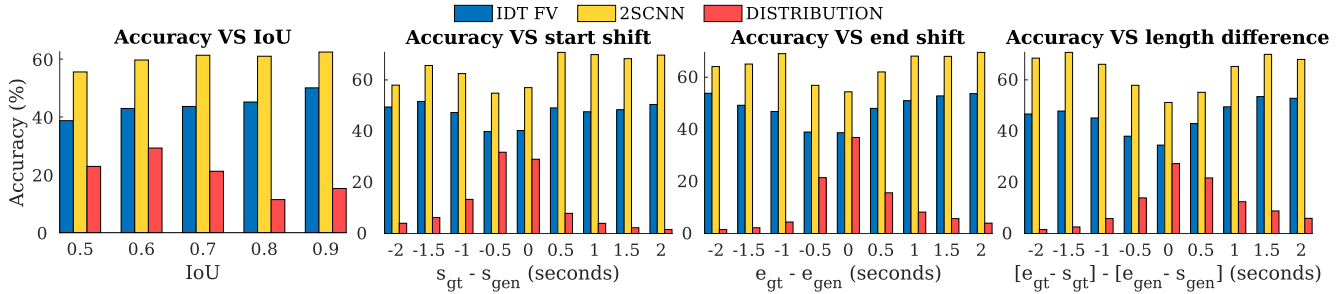


Figure 7: GTEA Gaze+: classification accuracy vs IoU, start/end shifts and length difference.

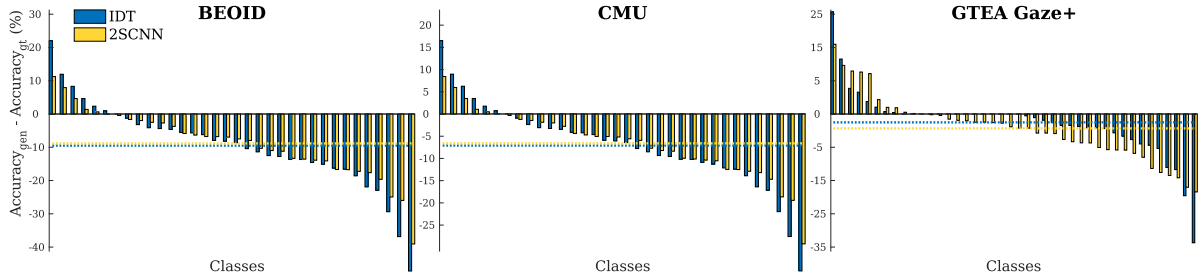


Figure 8: Accuracy per class differences. Most classes exhibit a drop in accuracy when testing generated segments.

and length differences), the generated segments achieve higher accuracy for both IDT FV and 2SCNN. When labels are inconsistent, shifting temporal bounds does not systematically alter the visual representation of the tested segments. The generated segments tend to include (or exclude) frames that increase the similarity between the testing and training segments.

Figure 8 reports per-class differences between generated and ground truth segments. Positive values entail that the accuracy for the given class is higher when testing the generated segments, and vice versa. Horizontal lines indicate the average accuracy difference. In total, 63% of classes in all three datasets exhibit a drop in accuracy drop when using IDT FV compared to 80% when using 2SCNN.

Dataset	2SCNN _{gt}	2SCNN _{gen}	2SCNN _{gt} ^{aug}	2SCNN _{gen} ^{aug}
BEOID	93.5	83.8	92.3	86.6
GTEA Gaze+	61.2	59.6	57.9	58.1

Table 3: 2SCNN data augmentation results.

Data augmentation: For completeness, we evaluate the performance when using temporal data augmentation methods on two datasets. Generated segments in Section 4.1 are used to augment training. We double the size of the training sets, taking random samples for augmentation. Test sets remained unvaried. Results are reported in Table 3. While we observe an increase in robustness, we also notice a drop in accuracy for ground truth segments, respectively of 1% and 4% for BEOID and GTEA Gaze+.

In conclusion, we note that both IDT FV and 2SCNN are sensitive to changes in temporal bounds for both consistent and inconsistent annotations. Approaches that improve robustness using data augmentation could be attempted, however a broader look at how the methods could be inherently more robust is needed, particularly for CNN architectures.

5. Labeling Proposal: The Rubicon Boundaries

The problem of defining consistent temporal bounds of an action is most akin to the problem of defining consistent bounding boxes of an object. Attempts to define guidelines for annotating objects' bounding boxes started nearly a decade ago. Among others, the VOC Challenge 2007 [4] proposed what has become the standard for defining the bounding box of an object in images. These consistent labels have been used to train state-of-the-art object detection and classification methods. With this same spirit, in this Section we propose an approach to consistently segment the temporal scope of an object interaction.

Defining RB: The Rubicon Model of Action Phases [10], developed in the field of Psychology, posits an action as a goal a subject desires to achieve and identifies the main sub-phases the person gets through in order to complete the action. First, a person decides what goal he wants to obtain. After forming his intention, he enters the so-called *pre-actional* phase, that is a phase where he plans to perform the action. Following this stage, the subject acts towards goal achievement in the *actional phase*. The two phases are delimited by three transition points: the initiation of prior motion, the start of the action and the goal achievement.




Figure 9: Rubicon Boundaries labeling examples for three object interactions.

The model is named after the historical fact of Caesar crossing the Rubicon river, which became a metaphor for deliberately proceeding past a point of no return, which in our case is the transition point that signals the beginning of an action. We take inspiration from this model, specifically from the aforementioned transitions points, and define two phases for an object interaction:

Pre-actional phase This sub-segment contains the preliminary motion that directly precedes the goal, and is required for its completion. When multiple motions can be identified, the pre-actional phase should contain only the last one; **Actional phase** This is the main sub-segment containing the motion strictly related to the fulfillment of the goal. The actional phase starts immediately after the pre-actional phase.

In the following section, we refer to a label as an RB annotation when the *beginning* of an object interaction aligns with the *start* of the pre-actional phase and the *ending* of the interactions aligns with the *end* of the actional phase.

Figure 9 depicts three object interactions labeled according to the Rubicon Boundaries. The top sequence illustrates the action of cutting a pepper. The sequence shows the subject fetching the knife before cutting the pepper and taking it off the plate. Based on the aforementioned definitions, the pre-actional phase is limited to the motion of moving the knife towards the pepper in order to slice it. This is directly followed by the actional phase where the user cuts the pepper. The actional phase ends as the goal of ‘cutting’ is completed. The middle sequence illustrates the action of opening a fridge, showing a person approaching the fridge, reaching towards the handle before pulling the fridge door open. In this case, the pre-actional phase would be the reaching motion, while the actional phase would be the pulling motion.

Evaluating RB: We evaluate our RB proposal for consistency, intuitiveness as well as accuracy and robustness.

(i) Consistency: We already reported consistency results in Section 3.2, where RB annotations exhibit higher average overlap and less variation for all the evaluated object interactions - average IoU for all pairs of annotators increased from 0.63 for conventional boundaries to 0.83 for RB. Figure 10 illustrates per-class IoU box plots for the pre-actional

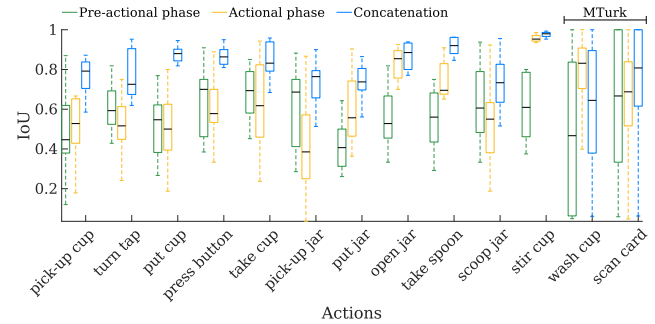


Figure 10: IoU comparison among the pre-actional phase (green), the actional phase (yellow) and their concatenation (blue) for several object interactions of BEOID.

and the actional phases separately, along with the concatenation of the two. For 7 out of the 13 actions, the actional phase was more consistent than the pre-actional phase, and for 12 out of the 13 actions, the concatenation of the phases proved the highest consistency.

(ii) Intuitiveness: While RB showed higher consistency in labeling, any new approach for temporal boundaries would require a shift in practice. We collect RB annotations from university students as well as from MTurk annotators. Locally, students successfully used the RB definitions to annotate videos with no assistance. However, this has not been the case for MTurk annotators for the two object interactions ‘wash cup’ and ‘scan card’. The MTurk HIT provided the formal definition of the *pre-actional* and *actional* phases, then ran two multiple-choice control questions to assess the ability of annotators to distinguish these phases from a video. The annotators had to select from textual descriptions what the pre and the actional phases entailed. For both object interactions, only a fourth of the annotators answered the control questions correctly.

Three possible explanations could be given, namely: annotators were accustomed to the conventional labeling method and did not spend sufficient time to study the definitions, or the definitions were difficult to understand. Further experimentation is needed to understand the cause.

(iii) Accuracy: We annotated GTEA Gaze+ using the Rubicon Boundaries, by employing three people to label its 1141 segments². For these experiments, we asked annotators to label both the pre-actional and the actional phase.

In Table 4, we report results for the actional phase alone (RB_{act}) as well as the concatenation of the two phases (RB), using 2SCNN on the same 5 folds from Section 4.2. The concatenated RB segments proved the most accurate, leading to an increase of more than 4% in accuracy compared to conventional ground truth segments. Temporal augmen-

²RB labels and video of results are available on project webpage: <http://www.cs.bris.ac.uk/~damen/Trespass/>

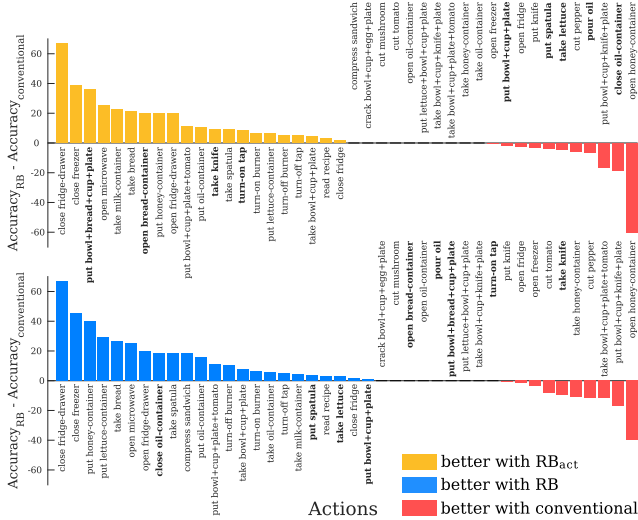


Figure 11: GTEA Gaze+: class accuracy difference between conventional and RB annotations. Some classes achieved higher accuracy only with RB_{act} , while other did only with the full RB segment. Bold highlights such cases.

tation on conventional labels ($Conv_{gt}^{aug}$) results in a drop of accuracy by 7.7% compared with the RB segments, highlighting that consistent labeling cannot be substituted with data augmentation. Figure 11 shows the accuracy per class difference between the two sets of RB annotations and the conventional labels. When using RB_{act} , 21/42 classes improved, whereas accuracy dropped for 11 classes compared to the conventional annotations. When using the full RB segment, 23/42 classes improved, while 10 classes were better recognized with the conventional annotations. In each case, 10 and 9 classes remain unchanged.

Given that the experimental setup was identical to that used for the conventional annotations, the boost in accuracy can be ascribed solely to the new action boundaries. Indeed, the RB approach helped the annotators to more consistently segment the object interactions contained in GTEA Gaze+, which is one of the most challenging datasets for egocentric action recognition.

(iv) Robustness: Table 4 also compares the newly annotated RB segments to generated segments with varied start and end times, as explained in Section 4.1. While RB_{gen} shows higher accuracy than the $Conventional_{gen}$ segments (59.6% as reported in Table 2), we still observe a clear drop in accuracy between gt and gen segments. Interestingly, we observe improved robustness when using the actional phase alone. Given that the actional segment’s start is closer in time to the beginning of the object interaction, when varying the start of the segment we are effectively including part of the pre-actional phase in the generated segment, which assists in making actions more discriminative.

Importantly, we show that RB annotations improved

$Conv_{gt}$	$Conv_{gt}^{aug}$	$act_{RB_{gt}}$	$act_{RB_{gen}}$	RB_{gt}	RB_{gen}
61.2	57.9	64.9	63.2	65.6	61.7

Table 4: GTEA Gaze+: 2SCNN classification accuracy comparison for conventional annotations (ground truth and augmented) and RB labels (ground truth and generated).

both consistency and accuracy of annotations on the largest dataset of egocentric object interactions. We believe these form solid basis for further discussions and experimentation on consistent labeling of temporal boundaries.

6. Conclusion and Future Directions

Annotating temporal bounds for object interactions is the base for supervised action recognition algorithms. In this work, we uncovered inconsistencies in temporal bound annotations within and across three egocentric datasets. We assessed the robustness of both hand-crafted features and fine-tuned end-to-end recognition methods, and demonstrated that both IDT FV and 2SCNN are susceptible to variations in start and end times. We then proposed an approach to consistently label temporal bounds for object interactions. We foresee three potential future directions:

Other NN architectures? While 2SCNN randomly samples frames from a video segment, the classification accuracy is still sensitive to variations in temporal bounds. Other architectures, particularly those that model temporal progression using Recurrent networks (including LSTM), rely on labeled training samples and would thus equally benefit from consistent labeling. Evaluating the robustness of such networks is an interesting future direction.

How can robustness to temporal boundaries be achieved? Classification methods that are inherently robust to temporal boundaries, while learning from supervised annotations, is a topic for future directions. As deep architectures reportedly outperform hand-crafted features and other classifiers, architectures that are designed to handle variations in start and end times are desired.

Which temporal granularity? The Rubicon Boundaries address consistent labeling of temporal bounds, but they do not address the concern of granularity of the action. Is the action of cutting a whole tomato composed of several short cuts or is it one long action? The Rubicon Boundaries model discusses actions relative to the goal a person wishes to accomplish. The granularity of an object interaction is another matter, and annotating the level of granularity consistently has not been addressed yet. Expanding Rubicon Boundaries to enable annotating the granularity would require further investigation.

Data Statement & Ack: Public datasets were used in this work; no new data were created as part of this study. RB annotations are available on the project’s webpage. Supported by EPSRC DTP and EPSRC LOCATE (EP/N033779/1).

References

- [1] D. Damen, T. Leelasawassuk, O. Haines, A. Calway, and W. Mayol-Cuevas. You-do, I-learn: Discovering task relevant objects and their modes of interaction from multi-user egocentric video. In *BMVC*, 2014.
- [2] F. De La Torre, J. Hodgins, A. Bargteil, X. Martin, J. Macey, A. Collado, and P. Beltran. Guide to the Carnegie Mellon University Multimodal Activity (CMU-MMAC) database. *Robotics Institute*, 2008.
- [3] O. Duchenne, I. Laptev, J. Sivic, F. Bach, and J. Ponce. Automatic annotation of human actions in video. In *ICCV*, 2009.
- [4] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (VOC) challenge. *IJCV*, 2010.
- [5] A. Fathi, A. Farhadi, and J. Rehg. Understanding egocentric activities. In *ICCV*, 2011.
- [6] A. Fathi, Y. Li, and J. Rehg. Learning to recognize daily actions using gaze. In *ECCV*, 2012.
- [7] A. Fathi, X. Ren, and J. Rehg. Learning to recognize objects in egocentric activities. In *CVPR*, 2011.
- [8] C. Feichtenhofer, A. Pinz, and A. Zisserman. Convolutional two-stream network fusion for video action recognition. In *CVPR*, 2016.
- [9] A. Gaidon, Z. Harchaoui, and C. Schmid. Temporal localization of actions with actoms. *TPAMI*, 2013.
- [10] P. M. Gollwitzer. Action phases and mind-sets. *Handbook of motivation and cognition: Foundations of social behavior*, 1990.
- [11] D. Huang, L. Fei-Fei, and J. C. Niebles. Connectionist temporal modeling for weakly supervised action labeling. In *ECCV*, 2016.
- [12] T. Kanade and M. Hebert. First-person vision. *Proceedings of the IEEE*, 100(8), 2012.
- [13] I. Laptev, M. Marszalek, C. Schmid, and B. Rozenfeld. Learning realistic human actions from movies. In *CVPR*, 2008.
- [14] C. Lea, A. Reiter, R. Vidal, and G. D. Hager. Segmental spatio-temporal cnns for fine-grained action segmentation and classification. In *ECCV*, 2016.
- [15] Y. J. Lee, J. Ghosh, and K. Grauman. Discovering important people and objects for egocentric video summarization. In *CVPR*, 2012.
- [16] Y. Li, Z. Ye, and J. Rehg. Delving into egocentric actions. In *CVPR*, 2015.
- [17] M. Ma, H. Fan, and K. Kitani. Going deeper into first-person activity recognition. In *CVPR*, 2016.
- [18] P. Matikainen, M. Hebert, and R. Sukthankar. Trajectons: Action recognition through the motion analysis of tracked features. In *ICCVW*, 2009.
- [19] W. Mayol, A. Davison, B. Tordoff, N. Molton, and D. Murray. Interaction between hand and wearable camera in 2D and 3D environments. In *BMVC*, 2004.
- [20] X. Peng and C. Schmid. Multi-region two-stream r-cnn for action detection. In *ECCV*, 2016.
- [21] H. Pirsiavash and D. Ramanan. Detecting activities of daily living in first-person camera views. In *CVPR*, 2012.
- [22] M. Ryoo, B. Rothrock, and L. Matthies. Pooled motion features for first-person videos. In *CVPR*, 2015.
- [23] J. Sánchez, F. Perronnin, T. Mensink, and J. Verbeek. Image classification with the fisher vector: Theory and practice. *IJCV*, 2013.
- [24] S. Satkin and M. Hebert. Modeling the temporal extent of actions. In *ECCV*, 2010.
- [25] K. Schindler and L. Van Gool. Action snippets: How many frames does human action recognition require? In *CVPR*, 2008.
- [26] G. A. Sigurdsson, G. Varol, X. Wang, A. Farhadi, I. Laptev, and A. Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *ECCV*, 2016.
- [27] S. Singh, C. Arora, and C. Jawahar. First person action recognition using deep learned descriptors. In *CVPR*, 2016.
- [28] K. Soomro, A. R. Zamir, and M. Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.
- [29] E. Spriggs, F. De La Torre, and M. Hebert. Temporal segmentation and activity classification from first-person sensing. In *CVPRW*, 2009.
- [30] O. Tange. GNU parallel - the command-line power tool. *The USENIX Magazine*, 36(1), Feb 2011.
- [31] E. Taralova, F. De La Torre, and M. Hebert. Source constrained clustering. In *ICCV*, 2011.
- [32] H. Wang and C. Schmid. Action recognition with improved trajectories. In *ICCV*, 2013.
- [33] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. Van Gool. Temporal segment networks: towards good practices for deep action recognition. In *ECCV*, 2016.
- [34] X. Wang, A. Farhadi, and A. Gupta. Actions transformations. In *CVPR*, 2016.
- [35] M. Wray, D. Moltisanti, W. Mayol-Cuevas, and D. Damen. SEMBED: Semantic embedding of egocentric action videos. In *ECCVW*, 2016.
- [36] X. Yang and Y. Tian. Effective 3D action recognition using eigenjoints. *Journal of Visual Communication and Image Representation*, 2014.
- [37] Y. Zhou and T. Berg. Temporal perception and prediction in ego-centric video. In *ICCV*, 2015.