

Therbligs in Action: Video Understanding through Motion Primitives

Eadom Dessalene, Michael Maynard, Cornelia Fermüller, Yiannis Aloimonos
University of Maryland, College Park
College Park, MD 20742, USA

{edessale,maynard,fermulcm,jyaloimo@umd.edu}

Abstract

In this paper we introduce a rule-based, compositional, and hierarchical modeling of action using Therbligs as our atoms. Introducing these atoms provides us with a consistent, expressive, contact-centered representation of action. Over the atoms we introduce a differentiable method of rule-based reasoning to regularize for logical consistency. Our approach is complementary to other approaches in that the Therblig-based representations produced by our architecture augment rather than replace existing architectures' representations. We release the first Therblig-centered annotations over two popular video datasets - EPIC Kitchens 100 and 50-Salads. We also broadly demonstrate benefits to adopting Therblig representations through evaluation on the following tasks: action segmentation, action anticipation, and action recognition - observing an average 10.5%/7.53%/6.5% relative improvement, respectively, over EPIC Kitchens and an average 8.9%/6.63%/4.8% relative improvement, respectively, over 50 Salads. Code and data will be made publicly available.

1. Introduction

We propose the use of *Therbligs* - a low-level mutually exclusive contact demarcated set of sub-actions. These Therbligs are **consistent** in that a given action segment has only a single Therblig representation, and Therbligs are **expressive** in that they capture the meaningful physical aspects of action relevant to action modeling. Therbligs were introduced in the early 20th century as a set of 18 elemental motions used to analyze complex movement - see the Supplementary Materials for a brief historical background. We adopt 7 Therbligs pertaining to those involving the manipulation of objects. See Figure 1 for our Therblig set.

The benefits of our Therblig-centered framework include: compositionality & hierarchy; rule-based reasoning; resolution of semantic ambiguity; contact-centered precision of temporal boundaries of action.

Contact transitions demarcate Therblig boundaries, giv-

ing Therbligs a consistency which methods relying on annotators' intuited demarcations lack ¹. Between points of contact exist contact states represented by a binary class (contact, no contact) for each object present, which are wholly captured by Therbligs. As objects in contact are the primary objects of interaction and define the space of possible actions, they provide meaningful information for the modeling of action.

Therblig atoms are then composable into higher entities, including full actions. These actions are in turn composable into sequences constituting activities. We then have the hierarchy of representation illustrated in Figure 2. At the lowest, and instantaneous, level are points of contact,

¹See [1] for a case study on how annotators have difficulties coming to a consensus on when actions begin and end.








Symbol	Name	Description
	Grasp (G)	When the worker's hand grabs the object
	Release (R)	The releasing of the object when it reaches its destination
	Hold (H)	The retention of an object such that it undergoes no movement while in operation
	Move (M)	Moving a loaded hand to the point of release, use, hold or (pre)position
	Reach (Re)	Moving an empty hand from the point of release
	Use (U)	When an object is being operated as intended
	Orient (O)	Changing the orientation of object while keeping it in roughly the same position

Figure 1. Listed above are the Therbligs we select, their symbolic illustrations as introduced by Gilbreths, and brief descriptions of their usage.

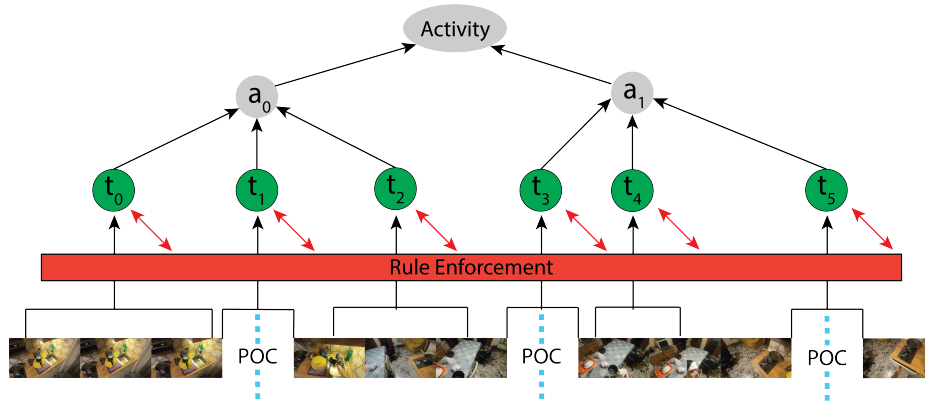


Figure 2. We introduce the use of **Therbligs** (t_i) in video understanding as a consistent, expressive, symbolic representation of sub-action. Points of Contact (indicated by the **divider dashes**) are necessarily associated with Therbligs and/or their boundaries. Because of the unambiguity of Points of Contact, Therblig boundaries gain precision and are non-overlapping. On top of Therblig atoms we construct a framework for **Rule Enforcement**, enforcing greater logical consistency through commonsense rules. This rule-based framework allows for the easy introduction of long-term constraints. Therblig atoms are then composable into actions (a_i), which are in turn composable into activities.

between which exist Therbligs with temporal extension, on top of which exist action, permutations of which constitute longer activities.

Architectures built upon Therbligs for the modeling of action gain temporal precision through points of contact as well as meaningful information captured by contact states. Therbligs also exhibit semantic mutual exclusivity in that there is one and only one Therblig interpretation of a sequence, as opposed to the many interpretations when action labels are intuited², leading to semantic ambiguity. As a consequence of Therbligs, semantic ambiguity at the action-level is constrained by the deeper grounding of action labels in explicit action dynamics (see Figure 2).

Unlike higher level actions, Therbligs enable the imposing of a contact-based logic defining preconditions and postconditions in the form of states of contact before and after Therbligs. For example, an object being *moved* must be preceded by *grasp* and proceeded by a *release*. The rules of this logic interface at the Therblig level of the hierarchy. These rules allow for biasing towards consistency between contact states and Therblig predictions within a loss term (see Section 3.2.2), and provide constraints over possible action sequences (see Section 3.2.1).

In producing sub-action level symbolic representations, our proposed hierarchical architecture is comprised of two main components; the **Therblig-Model**, which maps video to Therbligs; and, the **Action-Model**, which maps video and Therbligs to actions. The Therblig-Model is optimized over a loss including structure-aware terms for contact consistency and Therblig consistency by incorporating differentiable reasoning. Figure 3 illustrates our architecture.

²Some additional structure is needed for complete mutual exclusivity - see the Supplementary Materials for discussion on this structure.

This architecture is complementary to, rather than in competition with, existing architectures for action modeling - Therblig representations can be easily integrated through concatenation with existing feature representations. We demonstrate this with two state-of-the-art approaches to action segmentation - MSTCN++ [15] and ASFormer [22] and four popular approaches to action recognition - I3D [5], ViViT [2], TimeSFormer [3], and MoViNet [13].

We evaluate our approach over the tasks of action segmentation, action recognition, and action anticipation. We evaluate over the EPIC Kitchens 100 and 50-Salads datasets.

The primary contributions of our work are as follows:

- Therbligs, a consistent, expressive symbolic representation of sub-action centered on contact.
- Rules: Flexible and differentiable constraining of intuitive constraints on arrangement of atomic actions, informed by commonsense rules.
- Novel hierarchical architecture composed of a Therblig-Model and Action-Model. Representations produced by the Therblig-Model can be easily integrated into other approaches, as we demonstrate with six popular action understanding approaches.
- Dataset: We release the first Therblig-centered annotations over two popular video datasets.

The rest of this paper is structured as follows: Section 2 discusses related works, Section 3 introduces our proposed method, Section 4 describes the experiments, in Section 5 we provide discussion and in Section 6 we conclude.

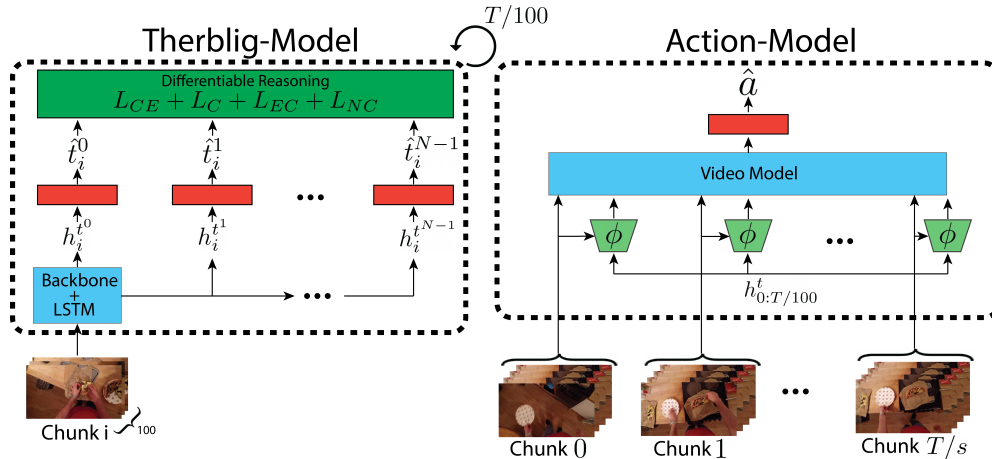


Figure 3. Architectural diagram of our framework. **Therblig-Model** takes a stack of $K = 100$ frames as input, feeding them to various backbone video architectures followed by a 2-layer GRU (**Backbone + LSTM**), which in turn produces hidden states h_i^t , passing through all $T/100$ stacks of the video. Hidden states h_i^t are fed to **fully connected** layers, followed by a Gumbel-Softmax operation, producing Therblig predictions amenable to differentiable reasoning. **Action-Model** takes a sliding window with window size W over the original video sequence with stride s , both values depending on the choice of architecture. These windows are fed to ϕ , an attention mechanism consisting of a 2 layer MLP - this MLP attends over the hidden states produced by **Therblig-Model**. The blended features produced by ϕ are fed together with the video window to a (**Video Network**) predicting action class likelihood a . See Sections 3.1.1 and 3.1.2 for details.

2. Related Works

2.1. Sub-Action Video Datasets

There exist several datasets that provide sub-action level annotations as a means of resolving semantic and temporal ambiguity in annotation, and enabling the hierarchical modeling of action [11, 19, 20]. FineGym [19] introduces fine-grained action annotations for actions in gymnastics, but suffers from a difficult and expensive data collection process. TAPOS [20] manually breaks actions into sub-actions for Olympics videos via temporal action parsing. Other datasets producing sub-action level annotations are of an instructional nature [14, 17, 21], providing annotations for steps of various cooking activities. Our Therblig annotations differ from other sub-action ontologies by 1) resolving temporal ambiguity by means of contact, 2) having a simple, logically consistent data collection process enabled through the imposing of commonsense rules, and 3) being flexible in application to a wide variety of datasets within the realm of object manipulation without relying on domain expertise.

2.2. Contact in Video Understanding

Contact has proven to be a useful feature in several tasks of interest in computer vision, such as hand-object pose estimation [4, 12, 18], character animation, kinematic pose estimation [16], etc. However, the vast majority of contact-centered approaches apply in the single-image setting, and few works [11] consider the modeling of contact for use in action understanding, despite contact being a defining char-

acteristic of all physical human interaction. Ego-OMG [7] approaches the task of action anticipation, representing long sequences of manipulation activity through sequences of discrete states, each state delineated by the making and breaking of contact. Rather than directly extract contact from each video frame as in [6, 7], we instead propose the adoption of Therbligs, a contact-centered representation governed by simple contact-based rules.

3. Methods

For video segment s_i with T frames, the Therblig representation t_i is a sequence of N Therblig tuples for every 100-frame chunk of s_i , and the contact representation c_i represents the objects in contact at the end of that video segment. Each Therblig annotation t_i is a sequence of the form $(v_0, o_0), \dots, (v_{N-1}, o_{N-1})$, where $v_j \in V$ and $V = \{\emptyset, Re, M, G, R, U, O, H\}$. In other words, v_j indicates the Therblig verb and each o_j indicates the object of interaction. See Figure 1 for full descriptions of each element of V (not included is \emptyset , corresponding to an empty sequence where no Therbligs occur). To illustrate, see the *Therbligs* and *Contacts* rows for annotations t_i and c_i , respectively, in Figure 6. We set the maximum number of possible Therblig annotations per sequence N to 6. Each contact annotation c_i is a tuple of the form (c_i^r, c_i^l) , where c_i^r corresponds to the class of the object held by the right hand, and c_i^l corresponds to the class of the object held by the left hand.

Given video segment s_i , our goal is to infer the relevant

action class likelihood over the tasks of action recognition and anticipation, and to infer the relevant action class likelihood for each frame over the task of action segmentation. We do this by means of a novel hierarchical architecture as described in Subsection 3.1. This architecture consists of two levels; a Therblig-Model (3.1.1) and an Action-Model (3.1.2). We then describe our rule-based reasoning formulation in Subsection 3.2, and detail the Therblig annotation collection process in Subsection 3.3.

3.1. Architecture

The architecture of our proposal is illustrated in Figure 3. We apply a backbone video network for each segment $s_i \in S$ where S is composed of video segments $S = \{s_0, \dots, s_{T/100}\}$, where T is the number of frames in S . This results in video features $F = \{f_0, \dots, f_{T/100}\}$. Our Therblig-Model predicts a sequence of Therbligs \hat{t}_i for each $f_i \in F$. Our Action-Model takes the representations produced by the Therblig-Model along with S , and produces action class likelihood(s) predictions.

As the Therblig annotations t_i and action annotations a_i do not exhibit one-to-one overlap between their respective video sequences, the Therblig-Model and Action-Model are trained separately.

3.1.1 Therblig-Model

The primary base architecture for Therblig-Model can be any popular video architecture. Features f_i are extracted from the backbone immediately prior to action classification and sent to a 2-layer GRU. Due to the lack of precise temporal alignment between the input video segments $s_i \in S$ and the Therblig annotations t_i , we adopt an encoder-decoder schema as follows: The hidden state of the GRU is set to f_i , and the network is rolled out to iteratively predict a sequence of Therbligs $\hat{t}_i = \{\hat{t}_i^0, \dots, \hat{t}_i^{N-1}\}$, feeding $\vec{0}$ as the initial input, and outputs of previous hidden layers as the inputs to the decoder for subsequent timesteps. We adopt the practice of teacher forcing, where the outputs of previous hidden layers are occasionally replaced with the ground truth, with probability $p = 0.5$. After training the Therblig-Model for 30 epochs and selecting the model instance with the top validation accuracy, we freeze the model for the training of the Action-Model.

3.1.2 Action-Model

As with the Therblig-Model, the base network of the Action-Model can be any popular video architecture. Sliding windows of size W is taken from input video S of T frames with a stride of s , where $s < W$ for the task of action segmentation, and $s = W$ for the task of action recognition & anticipation. Each window is fed to the video model, producing features $f_j \in \{f_0, \dots, f_{T/s}\}$ extracted immediately

prior to action classification. Our action segmentation models process the sliding windows sequentially, whereas our action recognition & anticipation base networks capture the entire video without windowing. We pair indices i from h_i^t and j from f_j by cross-referencing the closest times in video S associated with i and j .

We adopt a temporal attention mechanism ϕ , taking features f_j as input, and producing learned attention weights $\alpha_i^a = \{\alpha_i^{a_0}, \dots, \alpha_i^{a_{N-1}}\}$ over hidden layer outputs $\hat{h}_i = \{\hat{h}_i^{t_0}, \dots, \hat{h}_i^{t_{N-1}}\}$. After blending the hidden states with α_i^a , the blended features are concatenated with features f_j , and fed to fully connected layers to predict action class \hat{a} for action recognition & anticipation, and action classes $\hat{a}_j \in \{a_0, \dots, a_{T/s}\}$ for action segmentation.

3.2. Therblig Rules

Therbligs enable the introduction of contact-centered rules that 1) provide significant structure to the annotation process and 2) provide structure during training in the form of separate differentiable loss components. See Sections 3.2.1 and 3.2.2 for more.

3.2.1 Explicit Rules

Below we enumerate the explicit commonsense rules we introduce over the Therblig ontology.

Rule 1 The additions and subtractions of objects in contact (through grasping and releasing, respectively) produced by the Therblig sequence linking c_i and c_{i+N-1} must produce object contact set c_{i+N-1} from c_i .

Rule 2 Objects in contact state c_i cannot be *grasped* or *reached* without first being *released*.

Rule 3 Objects not in contact state c_i cannot be *moved*, *oriented*, *used*, or *released* without first being *grasped*.

These rules structure the second stage of crowdsourcing annotations in the form of a filter over all possible annotations. Annotations violating any of Rules 1, 2 or 3 are rejected and re-sent to annotators for correction - see Section 3.3 for details.

3.2.2 Differentiable Rules

We wish to incorporate the rules discussed in Section 3.2.1 into the training of the Therblig-Model. However, the rules are non-differentiable, and as such we approximate each rule with differentiable closed-form expressions. However, the logic requires discrete representations of Therblig predictions. As the softmax activation outputs are continuous, they must be converted into their discrete equivalents. The *argmax* operation is non-differentiable, and so we adopt the use of Gumbel Softmax as a differentiable alternative, applying it over the pre-softmax features of the Therblig-

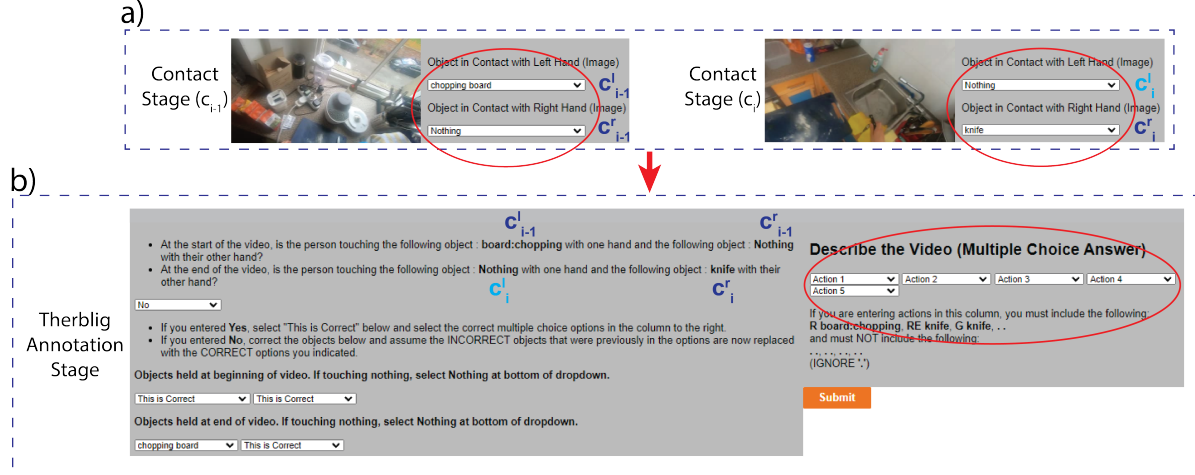


Figure 4. An overview of our two-stage crowdsourcing pipeline. The **Contact Annotation Stage** provides an image (shown) paired with a video (not shown) and asks the user to indicate the objects held by the actor in the image/video via multiple choice (circled). In the **Therblig Annotation Stage**, an annotator first validates the correctness of c_i and c_{i-1} and then produces Therblig annotations t_i . As can be seen above, in the Contact Annotation Stage a worker mistakenly indicated that nothing was held by the left hand for c_i . In the Therblig Annotation Stage, after correcting these erroneous contact annotations, a new worker annotates the sequence of Therbligs t_i (circled) - the multiple choice options shown to this worker have been filtered so as to achieve consistency with rules discussed in Section 3.2.1.

Model to arrive at discrete one-hot-encodings of Therbligs (\hat{g}_i) while maintaining differentiability.

The rules are represented as follows:

Rule 1 Cross Contact-Therblig Consistency Loss

$$L_C(i) = \sum_{k=0}^N \|c_i + \hat{g}_i^k \beta - c_{i+1}\| \quad (1)$$

Rule 2 Contact Enforcement Loss

$$L_{EC}(i) = \sum_{k=0}^N \|a_{i,k} - \hat{g}_i^k \gamma\| \quad (2)$$

where $a_{i,k} = a_{i,k-1} + \hat{g}_i^{k-1} \beta$ and $a_{i,0} = c_i$

Rule 3 Non-Contact Enforcement Loss

$$L_{NC}(i) = \sum_{k=0}^N \|a_{i,k} - \hat{g}_i^k \delta\| \quad (3)$$

where $a_{i,k} = a_{i,k-1} + \hat{g}_i^{k-1} \beta$ and $a_{i,0} = c_i$

Each of β , γ , and δ , correspond to a vector of size 7×1 , each index of which corresponds to a single Therblig verb. The predicted Therbligs \hat{g}_i^k are of dimension $|C| \times 7$, where C is the set of object categories in the dataset. Contact state c_i is of dimension $|C| \times 1$, taking values of 1 for object indices corresponding to those in contact with the hand, and 0 otherwise, based on ground truth. The k variable is used to iterate over each of the N elements in the Therblig sequence. The i variable takes a value in $[0, T/s]$ corresponding to segments in the input clip.

For β , values take on 1 for *grasp*, -1 for *release*, and 0 otherwise. These values of β were chosen such that $\hat{g}_i^k \beta$ reflects the addition and subtraction of objects to and from the hand. Our loss formulation in **Rule 1** penalizes predictions \hat{g}_i^k that do not produce contact state c_{i+1} from c_i .

For γ in **Rule 2**, values take on -1 for *reach* and *grasp*, 0 otherwise. These values of γ were chosen such that $\hat{g}_i^k \gamma$ results in loss that penalizes *reach* and *grasp* Therbligs when the Therblig-derived contact state $a_{i,k}$ already contains their respective objects.

For δ in **Rule 3**, values take on 1 for *move*, *orient*, *use*, or *release*, and 0 otherwise. These values of δ were chosen such that $\hat{g}_i^k \delta$ results in loss that penalizes those Therbligs when the Therblig-derived contact state $a_{i,k}$ does not contain their respective objects.

Loss component L_C of **Rule 1** measures the offset of c_i and c_{i+1} against $\hat{g}_i^k \beta$. Loss component L_{EC} of **Rule 2** iteratively compares the Therblig-derived contact state $a_{i,k}$ against $\hat{g}_i^k \gamma$. Loss component L_{NC} of **Rule 3** iteratively compares $a_{i,k}$ against $\hat{g}_i^k \delta$.

We adopt each of these loss terms in addition to Categorical Cross Entropy loss L_{CE} to arrive at combined loss $L = L_{CE} + L_{NC} + L_{EC} + L_C$. These loss terms provide meaningful constraints in guiding the learning process, and we demonstrate this finding empirically in Section 4.2.

3.3. Crowdsourcing

We crowdsource our annotations in two stages; see Figure 4. The first stage involves the crowdsourcing of objects in contact with the hands. The annotator is shown an im-

age, along with a slightly slowed-down video roughly 5 seconds long. The addition of the video adds temporal context for images too difficult to annotate alone. The annotator is asked to indicate via multiple choice the objects in contact with the actor’s hands at the end of the interaction. We pool 5 independently collected responses per image, taking the mode of the answers for consistency. The average time per assignment was 3.4 minutes.

In the second stage, the possible Therbligs for each item in the sequence t_i^j for $0 \leq j < N$ are determined by taking the cross product of the contact annotations c_i and c_{i+1} with $V = \{Re, M, G, R, O, U, H\}$ and filter the results such that *consistency* is observed with respect to the three rules defined in Subsection 3.2.1. **We note there are 1,067 possible Therblig (verb, object) tuples; through our rules, we are able to significantly reduce the average number of multiple choice annotation possibilities to just 19 (verb, object) tuples!** We set N , the number of Therblig tuples per video clip, to 6.

4. Experiments

Our experiments on action segmentation explore the extent to which we are able to predict Therbligs and the extent to which their incorporation benefits action segmentation³. We report our numbers by early stopping over validation accuracy for 3 independent runs, reporting only the mean. All code, data and results will be released upon acceptance.

4.1. Datasets

EPIC Kitchens We choose the EPIC Kitchens 100 dataset because of the benefits the egocentric perspective provides in allowing for full view of the hands and objects in contact. We augment portions of the EPIC Kitchens dataset

³See [here](#) for an exhaustive list of videos paired with their corresponding Therblig annotations, contact state annotations, Therblig predictions and action annotations.

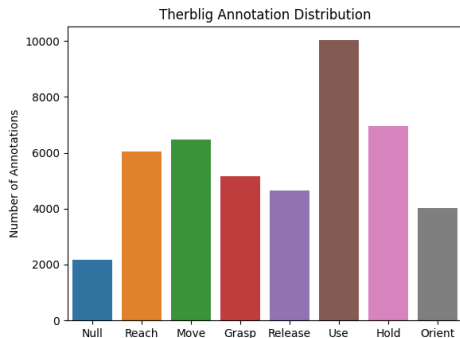


Figure 5. Breakdown by frequency of Therblig verbs over all collected annotations. Null corresponds to Therblig sequences where workers indicated no hand-object activity took place.

Full Data	50-Salads	EPIC Kitchens
L_{CE}	22.1%/3.19/2.25	9.5%/5.68/2.491
$L_{CE} + L_C$	23.0%/3.15/1.96	9.9%/5.209/2.124
$L_{CE} + L_{EC}$	23.2%/2.99/1.92	11.7%/5.238/1.90
$L_{CE} + L_{NC}$	23.1%/2.83/1.98	12.7%/5.16/1.91
All L	25.1%/2.6/1.70	13.7%/5.19/1.83
Low Data		
L_{CE}	12.1%/5.91/2.53	7.4%/5.37/2.36
All L	16.9%/5.25/1.96	10.1%/5.21/1.86

Table 1. Evaluation of Therblig-Model when trained over all Therblig annotations (Full Data) and when trained over a subset (Low Data). Results reported in order of : Accuracy \uparrow /Levenshtein Distance \downarrow / Logical Consistency \downarrow .

with densely labeled Therbligs, for a total of 14,600 crowdsourced annotations. In accordance with reporting practices of other action segmentation approaches over EPIC Kitchens, we perform training and evaluation solely at the verb-level. See the Supplementary Materials for more.

50 Salads We choose the 50 Salads dataset because each activity performed is strongly structured by the making and breaking of contact. We augment the entirety of the video in this dataset, for a total of 6,500 crowdsourced annotations. See the Supplementary Materials for more.

4.2. Therblig Prediction

In this set of experiments, we answer the extent to which our Therblig-Model is capable of mapping video chunks s_i to the sequence of Therbligs $\hat{t}_i = \{\hat{t}_i^0, \dots, \hat{t}_i^{N-1}\}$.

Metrics We evaluate our results, comparing \hat{t}_i^j and t_i^j , for $0 \leq j \leq N, \forall i$ over the following metrics of evaluation: Element-wise accuracy and Levenshtein distance (**L**). Element-wise accuracy metrics suffer from the strict ordering requirement; and so are not reflective of sequence-level similarity. Therefore we also evaluate over Levenshtein distance, the number of edits (insertions, deletions, and swaps) to transform \hat{t}_i^j into t_i^j . In addition, we evaluate the logical consistency of our predictions as measured by the normalized number of violations per sequence of the rules described in Section 3.2.1.

Comparisons Table 1 illustrates the results of various forms of Therblig-Model using a backbone I3D network over the EPIC Kitchens and 50 Salads datasets. L_{CE} refers to a simple, 2-layer bidirectional GRU trained solely over Categorical Cross Entropy. We train our GRU with and without each of the loss components discussed in Section 3.2.2. In addition, we include results when training over $L_{CE} + L_C + L_{EC} + L_{NC}$ and L_{CE} , but in the low-data setting where roughly 10% of our annotations are trained



Figure 6. Qualitative results for the action segmentation task showing Therblig annotations (*Therbligs*), contact annotations c_i (*contacts*), ground truth annotations (*GT*), the base GRU baseline (*GRU*) and the base GRU with Therbligs (*GRU (T)*). Grey colors indicate unannotated frames. Errors in *GRU* due to lack of mutual exclusivity are shown in the form of non-black contrastive colors w.r.t. *GT*. Black colors indicate all other action misclassifications. Actual model predictions are at the verb-level (objects are introduced into predictions solely for purposes of illustration).

	50 Salads	EPIC Kitchens
GRU	62.8%/8.45/60.9	39.1%/19.3/23.2
GRU w. T	66.9%/8.21/65.1	44.7%/13.5/25.6
Base MST	78.19%/7.1/74.16	53.44%/10.4/54.38
MST w. T	85.25%/6.7/81.01	58.75%/9.9/60.97
Base ASF	82.97%/6.9/77.23	58.14%/10.3/56.41
ASF w. T	88.2%/6.3/84.33	64.03%/9.5/61.46

Table 2. Action Segmentation results over EPIC Kitchens and 50 Salads datasets for frame-wise accuracy \uparrow /segmental edit distance \downarrow and segmental F1-score@25 \uparrow . Base models (without **w. T**) are followed by base models incorporating Therbligs (**w. T**).

over for both datasets, highlighting the value of the structure defined in Section 3.2.2. See Table 1 for results. All training and evaluation details are described in the Supplementary Materials section.

4.3. Action Segmentation

In this experiment we evaluate the extent to which the incorporation of Therblig-Model benefits performance in action segmentation. In accordance with [8] we perform predictions at the verb level, rather than the verb-object level.

Metrics As action segmentation is a classic task in computer vision, we rely on the works of [8, 10] in adopting the following evaluation metrics: frame-wise accuracy, segmental edit-score, and segmental F1-score. Frame-wise accuracy is used in all action segmentation works, whereas segmental edit-score and F1-score are most commonly used to penalize over-segmentation in particular.

Comparisons Table 2 illustrates the results of ablations of our proposed architecture over the EPIC Kitchens and 50 Salads datasets for the following base architectures of Action-Model: **GRU**, **MST** (MSTCN++), and **ASF** (ASFormer). The models (with no suffix) correspond to the mapping of raw video to framewise action labels and the models with the (**w. T**) suffix correspond to those of our proposed framework. For simplicity, we utilize the same Therblig-Model across all models with an I3D backbone, as ASFormer and MSTCN++ both use I3D features as their respective backbones. All training and evaluation details are described in the Supplementary Materials section.

4.4. Action Recognition & Anticipation

In these experiments we evaluate the extent to which the incorporation of Therblig-Model benefits performance in action recognition & anticipation. The Therblig-model and Action-model backbones are identical throughout all these experiments.

Metrics We adopt accuracy as the primary metric of comparison for all action recognition & anticipation results. We structure results over EPIC Kitchens by verb recognition accuracies, object recognition accuracies, and action recognition accuracies.

Comparisons Tables 3 and 4 illustrate the results of ablations of our proposed architecture over the EPIC Kitchens and 50 Salads datasets for the following base architectures of Action Model: **I3D**, **TimeSFormer** (TimeSFormer-B), **ViViT** (ViViT-B/16x2) and **MoViNet** (MoViNet-A3). The models (with no suffix) correspond to the mapping of raw video to framewise action labels and the models with the (**w. T**) suffix correspond to those of our

	50 Salads	EPIC Kitchens
I3D	64.2%	65.1%/49.7%/39.2%
I3D w. A	64.0%	65.9%/49.6%/39.5%
I3D w. T	69.1%	68.2%/51.9%/41.6%
TimeSFrmr	73.1%	62.1%/55.4%/41.1%
TimeSFrmr w. A	74.2%	61.9%/56.9%/41.9%
TimeSFrmr w. T	76.4%	67.6%/57.0%/44.0%
ViViT	72.1%	62.4%/56.0%/43.3%
ViViT w. A	72.9%	64.0%/57.2%/44.0%
ViViT w. T	74.3%	66.4%/56.9%/45.9%
MoViNet	73.9%	67.9%/52.9%/41.1%
MoViNet w. A	72.8%	67.6%/52.1%/40.2%
MoViNet w. T	76.5%	69.1%/55.0%/43.8%

Table 3. Action recognition accuracies over EPIC Kitchens and 50 Salads datasets. Results under EPIC Kitchens are provided as: verb/object/action prediction accuracies, respectively. Base models (without **w. T**) are followed by base models incorporating action labels (**w. A**) and Therbligs (**w. T**).

proposed framework. All models are pre-trained over Kinetics 400. For purposes of reproducibility, we describe all details in the Supplementary Materials.

5. Discussion

To give better intuition we show predicted Therblig sequences alongside video of manipulation activity [here](#).

We point the reader’s attention towards the ”Low Data” results reported in Table 1, where we observe a large increase in accuracy through the incorporation of all the rule-based loss components. We believe this validates our hypothesis that the constraints imposed by the rules play a particularly outsized role in the low-data setting, where the model would otherwise have to infer the same common-sense structure, motivating future possible directions in the few-shot domain.

As demonstrated in Tables 2, 3, and 4, the incorporation of Therbligs results in superior performance over both the 50 Salads and EPIC Kitchens 100 datasets for all base architectures and tasks. While the accuracies of the models trained in Figure 2 come close to matching the reported numbers in the original papers, the accuracy of the models trained with Therbligs reported in Table 2 outperform baseline models trained by us as well as those reported in original papers.

We point the reader’s attention to the particularly pronounced role Therbligs play in improving verb recognition accuracy in the action segmentation (Table 2) and action recognition (Table 3) results, even more pronounced among the Transformer models, likely due to difficulties of

	50 Salads	EPIC Kitchens
I3D	39.5%	31.1%/18.3%/10.1%
I3D w. A	41.9%	32.0%/18.1%/10.4%
I3D w. T	43.1%	33.0%/18.9%/10.9%
TimeSFrmr	48.6%	31.6%/28.2%/13.6%
TimeSFrmr w. A	49.0%	29.9%/28.5%/13.1%
TimeSFrmr w. T	51.4%	33.9%/28.5%/14.8%
ViViT	45.6%	31.9%/29.8%/13.9%
ViViT w. A	44.9%	33.0%/ 30.1% /14.5%
ViViT w. T	49.1%	33.7%/29.4%/14.7%
MoViNet	46.4%	34.2%/25.6%/13.1%
MoViNet w. A	47.2%	34.9%/26.0%/13.3%
MoViNet w. T	48.2%	36.1%/26.9%/14.0%

Table 4. Action anticipation accuracies over EPIC Kitchens and 50 Salads datasets. Results under EPIC Kitchens are provided as: verb/object/action prediction accuracies, respectively. Base models (without **w. T**) are followed by base models incorporating action labels (**w. A**) and Therbligs (**w. T**).

Transformers in capturing the fine-grained aspects of motion [9]. We also point the reader to Figure 6 for qualitative results. When more than one action applies to any given video segment, the model incorporating Therbligs is more likely to produce interpretations of action that happen to align more closely with actions belonging in the ground truth. We hypothesize the reason to be that the incorporation of Therbligs alleviates the burden on the network of learning movement representations, and allows it to focus on contextual cues important to resolving issues of mutual exclusivity. Furthermore, the model incorporating Therbligs produces predictions with better action boundaries.

We point the reader to the Supplementary Material for more discussion.

6. Conclusion

In this paper we have presented a method for mitigation of common limitations in action understanding approaches using a novel framework structured around Therbligs - a consistent, expressive, contact-centered representation of action. We demonstrate superior performance through the introduction of Therbligs over action recognition/action anticipation/action segmentation - observing an average 10.5%/7.53%/6.5% improvement, respectively, over EPIC Kitchens and an average 8.9%/6.63%/4.8% relative improvement, respectively, over 50 Salads. We hope the release of Therblig annotations inspires future work towards the hierarchical modeling of action, and will release all code and data upon acceptance.

References

- [1] Humam Alwassel, Fabian Caba Heilbron, Victor Escorcia, and Bernard Ghanem. Diagnosing error in temporal action detectors. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 256–272, 2018.
- [2] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6836–6846, 2021.
- [3] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, volume 2, page 4, 2021.
- [4] Zhe Cao, Ilija Radosavovic, Angjoo Kanazawa, and Jitendra Malik. Reconstructing hand-object interactions in the wild. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12417–12426, 2021.
- [5] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017.
- [6] Eadom Dessalene, Chinmaya Devaraj, Michael Maynard, Cornelia Fermuller, and Yiannis Aloimonos. Forecasting action through contact representations from first person video. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [7] Eadom Dessalene, Michael Maynard, Chinmaya Devaraj, Cornelia Fermuller, and Yiannis Aloimonos. Egocentric object manipulation graphs. *arXiv preprint arXiv:2006.03201*, 2020.
- [8] Yifei Huang, Yusuke Sugano, and Yoichi Sato. Improving action segmentation via graph-based temporal reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14024–14034, 2020.
- [9] Ziyuan Huang, Zhiwu Qing, Xiang Wang, Yutong Feng, Shiwei Zhang, Jianwen Jiang, Zhurong Xia, Mingqian Tang, Nong Sang, and Marcelo H Ang Jr. Towards training stronger video vision transformers for epic-kitchens-100 action recognition. *arXiv preprint arXiv:2106.05058*, 2021.
- [10] Yuchi Ishikawa, Seito Kasai, Yoshimitsu Aoki, and Hirokatsu Kataoka. Alleviating over-segmentation errors by detecting action boundaries. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2322–2331, 2021.
- [11] Jingwei Ji, Ranjay Krishna, Li Fei-Fei, and Juan Carlos Niebles. Action genome: Actions as compositions of spatio-temporal scene graphs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10236–10247, 2020.
- [12] Korrawe Karunratanakul, Jinlong Yang, Yan Zhang, Michael J Black, Krikamol Muandet, and Siyu Tang. Grasping field: Learning implicit representations for human grasps. In *2020 International Conference on 3D Vision (3DV)*, pages 333–344. IEEE, 2020.
- [13] Dan Kondratyuk, Liangzhe Yuan, Yandong Li, Li Zhang, Mingxing Tan, Matthew Brown, and Boqing Gong. Movinets: Mobile video networks for efficient video recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16020–16030, 2021.
- [14] 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 conference on computer vision and pattern recognition*, pages 780–787, 2014.
- [15] 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*, 2020.
- [16] Davis Remppe, Leonidas J Guibas, Aaron Hertzmann, Bryan Russell, Ruben Villegas, and Jimei Yang. Contact and human dynamics from monocular video. In *European Conference on Computer Vision*, pages 71–87. Springer, 2020.
- [17] Marcus Rohrbach, Anna Rohrbach, Michaela Regneri, Sikandar Amin, Mykhaylo Andriluka, Manfred Pinkal, and Bernt Schiele. Recognizing fine-grained and composite activities using hand-centric features and script data. *International Journal of Computer Vision*, pages 1–28, 2015.
- [18] Dandan Shan, Jiaqi Geng, Michelle Shu, and David F Fouhey. Understanding human hands in contact at internet scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9869–9878, 2020.
- [19] Dian Shao, Yue Zhao, Bo Dai, and Dahua Lin. Finegym: A hierarchical video dataset for fine-grained action understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2616–2625, 2020.
- [20] Dian Shao, Yue Zhao, Bo Dai, and Dahua Lin. Intra-and inter-action understanding via temporal action parsing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 730–739, 2020.
- [21] Sebastian Stein and Stephen J McKenna. Combining embedded accelerometers with computer vision for recognizing food preparation activities. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, pages 729–738, 2013.
- [22] Fangqiu Yi, Hongyu Wen, and Tingting Jiang. As-former: Transformer for action segmentation. *arXiv preprint arXiv:2110.08568*, 2021.