

Every Moment Counts: Dense Detailed Labeling of Actions in Complex Videos

Serena Yeung
Stanford University
syyeung@cs.stanford.edu

Olga Russakovsky
Stanford University
olga@cs.stanford.edu

Ning Jin
Stanford University
jennyjn@cs.stanford.edu

Mykhaylo Andriluka
Stanford University
andruluk@stanford.edu

Greg Mori
Simon Fraser University
mori@cs.sfu.ca

Li Fei-Fei
Stanford University
feifeili@cs.stanford.edu

Abstract

Every moment counts in action recognition. A comprehensive understanding of human activity in video requires labeling every frame according to the actions occurring, placing multiple labels densely over a video sequence. To study this problem we extend the existing THUMOS dataset and introduce MultiTHUMOS, a new dataset of dense labels over unconstrained internet videos. Modeling multiple, dense labels benefits from temporal relations within and across classes. We define a novel variant of long short-term memory (LSTM) deep networks for modeling these temporal relations via multiple input and output connections. We show that this model improves action labeling accuracy and further enables deeper understanding tasks ranging from structured retrieval to action prediction.

1. Introduction

Humans are great at multi-tasking: they can be walking while talking on the phone while holding a cup of coffee. Further, human action is continual, and every minute is filled with potential labeled action (Fig. 1). However, most work on human action recognition in video focuses on recognizing discrete instances or single actions at a time: for example, which sport [8] or which single cooking activity [21] is taking place. We argue this setup is fundamentally limiting. First, a single description is often insufficient to fully describe a person’s activity. Second, operating in a single-action regime largely ignores the intuition that actions are intricately connected. A person that is running and then jumping is likely to be simultaneously doing a sport such as basketball or long jump; a nurse that is taking a patient’s blood pressure and looking worried is likely to call a doctor as her next action. In this work, we go beyond the standard one-label paradigm to dense, detailed, multilabel understanding of human actions in videos.




Figure 1: In most internet videos there are multiple simultaneous human actions. Here, we show a concrete example from a basketball video to illustrate our target problem of dense detailed multi-label action understanding.

There are two key steps on the path to tackling detailed multilabel human action understanding: (1) finding the right dataset and (2) developing an appropriate model. In this paper we present work in both dimensions.

The desiderata for a video dataset include the following: video clips need to be long enough to capture multiple consecutive actions, multiple simultaneous actions need to be annotated, and labeling must be dense with thorough coverage of action extents. Video annotation is very time-consuming and expensive, and to the best of our knowledge no such dataset currently exists. UCF101 [28], HMDB51 [12], and Sports1M [8] are common challenging action recognition datasets. However, each video is associated with non-localized labels (Sports1M), and the videos in UCF101 and HMDB51 are further temporally clipped around the action. MPII Cooking [21] and Breakfast [11] datasets contain long untrimmed video sequences with multiple sequential actions but still only one label per frame; further, they are restricted to closed-world kitchen environments. THUMOS [7] contains long untrimmed videos but

most videos (85%) only contain a single action class.

To overcome these problems, we introduce a new action detection dataset called MultiTHUMOS, significantly extending the annotations on 413 videos (30 hours) of THUMOS action detection dataset. First, MultiTHUMOS allows for an in-depth study of simultaneous human action in video: it extends THUMOS from 20 action classes with 0.3 labels per frame to 65 classes and 1.5 labels per frame. Second, MultiTHUMOS allows for a thorough study of the temporal interaction between consecutive actions: the average number of distinct action categories in a video is 10.5 (compared to 1.1 in THUMOS). Going further, MultiTHUMOS lends itself to studying intricate relationships between action labels: the 45 new annotated classes include relationships such as hierarchical (e.g., more general Throw or PoleVault and more specific BasketballShot or PoleVaultPlantPole) and fine-grained (e.g., Guard versus Block or Dribble versus Pass in basketball). Fig. 1 shows an example of our dense multilabel annotation.

Reasoning about multiple, dense labels on video requires models capable of incorporating temporal dependencies. Recent action recognition literature has used recurrent neural networks known as Long Short Term Memory (LSTM) for action recognition in videos [3]. We introduce MultiLSTM, a new LSTM-based model targeting dense, multilabel action analysis. Taking advantage of the fact that more than 45% of frames in MultiTHUMOS have 2 or more labels, the model can learn dependencies between actions in nearby frames and between actions in the same frame, which allows it to subsequently perform dense multilabel temporal action detection on unseen videos.

In summary, our contributions are:

1. We introduce MultiTHUMOS, a new large-scale dataset of dense, multilabel action annotations in temporally untrimmed videos, and
2. We introduce MultiLSTM, a new recurrent model based on an LSTM that features temporally-extended input and output connections.

Our experiments demonstrate improved performance of MultiLSTM relative to a plain LSTM baseline on our dense, multilabel action detection benchmark.

2. Related Work

Visual analysis of human activity has a long history in computer vision research. Thorough surveys of the literature include Poppe [20] and Weinland et al. [36]. Here we review recent work relevant to dense labeling of videos.

Datasets. Research focus is closely intertwined with dataset creation and availability. The KTH dataset [24] was a catalyst for a body of work. This era focused on recognizing individual human actions, based on datasets con-

	Detection	Untrimmed	Open-world	Multilabel
UCF101 [28]	-	-	yes	-
HMDB51 [12]	-	-	yes	-
Sports 1M [8]	-	yes	yes	-
Cooking [21]	yes	yes	-	-
Breakfast [11]	yes	yes	-	-
THUMOS [7]	yes	yes	yes	-
MultiTHUMOS	yes	yes	yes	yes

Table 1: Our MultiTHUMOS dataset overcomes many limitations of previous datasets.

sisting of an individual human imaged against a generally stationary background. In subsequent years, the attention of the community moved towards more challenging tasks, especially in the level of background clutter present in the videos. The HMDB [12], UCF 101 [28], and THUMOS [7] datasets exemplify this trend toward utilizing “unconstrained” internet video as a dataset source. Task direction has also moved toward a retrieval setting, finding a (small) set of videos from a large background collection, including datasets such as TRECVID MED [18] and Sports 1M [8].

While the push toward unconstrained internet video is positive in terms of the difficulty of this task, it has moved focus away from human action toward identifying scene context. Discriminating diving versus gymnastics largely involves determining the scene of the event. The MPII Cooking dataset [21] and Breakfast dataset [11] refocus efforts toward human action within restricted action domains (Table 1). The MultiTHUMOS dataset we propose shares commonalities with this line, but emphasizes generality of video, multiple labels per frame, and a broad set of general to specific actions.

Deep learning for video. In common with object recognition, hand-crafted features for video analysis are giving way to deep convolutional feature learning strategies. The best hand-crafted features, the dense trajectories of Wang et al. [34], achieve excellent results on benchmark action recognition datasets. However, recent work has shown superior results by learning video features (often combined with dense trajectories). Simonyan and Zisserman [26] present a two-stream convolutional architecture utilizing both image and optical flow data as input sources. Zha et al. [40] examine aggregation strategies for combining deep learned image-based features for each frame, obtaining impressive results on TRECVID MED retrieval. Karpathy et al. [8] and Tran et al. [32] learn spatio-temporal filters in a deep network and apply them to a variety of human action understanding tasks. Mansimov et al. [14] consider methods for incorporating ImageNet training data to assist in initializing model parameters for learning spatio-temporal features. Wang et al. [35] study temporal pooling strategies, specifically focused on classification in variable-length input videos.

Temporal models for video. Constructing models of the temporal evolution of actions has deep roots in the literature. Early work includes Yamato et al. [38], using hidden Markov models for latent action state spaces. Discriminative variants include those based on latent SVMs over key poses and action grammars [17, 33, 19]. A recent set of technical reports has deployed deep models using long short-term memory (LSTM) models [5] for video analysis [3, 15, 29, 39]. These papers have shown promising results applying LSTMs for tasks including video classification and sentence generation. In contrast, we develop a novel LSTM that performs spatial input aggregation and output modeling for dense labeling output.

Action detection. Beyond assigning a single label to a whole video, the task of action detection localizes this action within the video sequence. An example of canonical work in this vein is Ke et al. [9]. More recent work extended latent SVMs to spatio-temporal action detection and localization [30, 13]. Rohrbach et al. [22] detect cooking actions using hand-centric features accounting for human pose variation. Ni et al. [16] similarly utilize hand-centric features on the MPII Cooking dataset, but focus on multiple levels of action granularity. Gkioxari and Malik [4] train SVMs for actions on top of deep learned features, and further link them in time for spatio-temporal action detection. In contrast, we address the task of dense multilabel action detection.

Attention-based models. Seminal work on computational spatial attention models for images was done by Itti et al. [6]. Recent action analysis work utilizing attention includes Shapovalova et al. [25] who use eye-gaze data to drive action detection and localization. Xu et al. [37] use visual attention to assist in caption generation. Yao et al. [39] develop an LSTM for video caption generation with soft temporal attention. Our method builds on these directions, using an attention-based input temporal context for dense action labeling.

3. The MultiTHUMOS Dataset

Research on detailed, multilabel action understanding requires a dataset of untrimmed, densely labeled videos. However, we are not aware of any existing dataset that fits these requirements. THUMOS [7] is untrimmed but contains on average only a single distinct action labeled per video. MPII Cooking [21] and Breakfast [11] datasets have labels of sequential actions, but contain only a single label per frame and are further captured in closed-world settings of a single or small set of kitchens (Table 1).

To address the limitations of previous datasets, we introduce a new dataset called MultiTHUMOS. MultiTHUMOS contains dense, multilabel, frame-level action annotations (Fig. 2) for 30 hours across 400 videos in the THUMOS

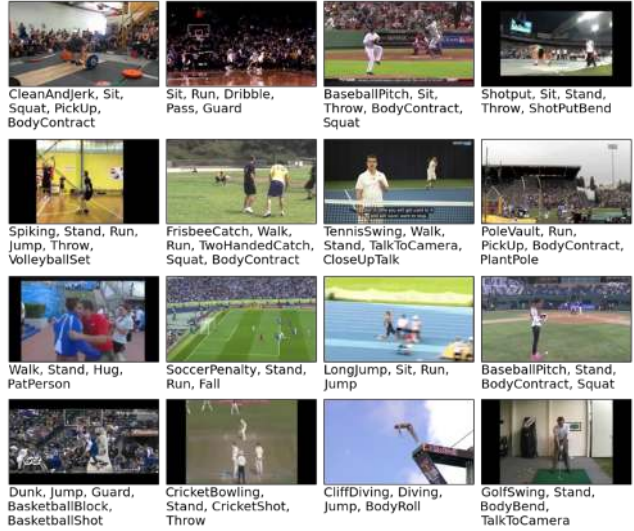


Figure 2: Our MultiTHUMOS dataset contains multiple action annotations per frame.




Figure 3: *Left.* MultiTHUMOS has significantly more labels per frame than THUMOS [7] (1.5 in MultiTHUMOS versus 0.3 in THUMOS). *Right.* Additionally, MultiTHUMOS contains up to 25 action labels per video compared to ≤ 3 labels in THUMOS.

’14 action detection dataset. In particular, all videos in the “Validation Data” and “Test Data” sets were labeled.¹ Annotations were collected in collaboration with a commercial crowdsourcing service. Workers were provided with the name of an action, a brief (up to 1 sentence) description, and 2 annotation examples, and asked to annotate the start and end frame of the action in the videos. An action was annotated if it occurred anywhere in the frame.

In total, we collected 32,325 annotations of 45 action classes, bringing the total number of annotations from 6,365 over 20 classes in THUMOS to 38,690 over 65 classes in MultiTHUMOS. The density of annotations increased from 0.3 to 1.5 labels per frame on average (with the maximum number per frame increasing from 2 in THU-

¹ THUMOS training data consists of 3 sets of videos: temporally clipped “Training Data”, temporally untrimmed “Validation Data” with temporal annotations, and temporally untrimmed “Background Data” with no temporal annotations. Test data consists of temporally untrimmed “Test Data” with temporal annotations. We annotated all videos sets originally including temporal annotations, i.e. “Validation Data” and “Test Data”.



Figure 5: Our MultiTHUMOS dataset is very challenging due to high intra-class variation.

MOS to 9 in MultiTHUMOS), and from 1.1 to 10.5 action classes per video. Maximum number of actions per video increased from 3 in THUMOS to 25 in MultiTHUMOS. Fig. 3 shows the full distribution of annotation density.

Using these dense multilabel video annotations, we are able to learn and visualize the relationships between actions. The co-occurrence hierarchy of object classes in images based on mutual information of object annotations was learned by Choi et al. [2]; we adapt their method to per-frame action annotations in video. Fig. 4 shows the resulting action hierarchy. Classes such as Squat and BodyContract frequently co-occur; in contrast, classes such as Run and Billiards rarely occur together in the same frame.

MultiTHUMOS is a very challenging dataset for three key reasons. First, it has many fine-grained action categories with low visual inter-class variation, including hierarchical (e.g. throw vs. baseball pitch), hierarchical within a sport (e.g. pole vault vs. the act of planting the pole when pole vaulting), and fine-grained (e.g. basketball dunk, shot, dribble, guard, block, and pass). We also chose both sport-specific actions (such as different basketball or volleyball moves), as well as general actions that can occur in multiple sports (e.g. pump fist, or one-handed catch). Second, it has high intra-class variation as shown in Fig. 5. The same



Figure 6: MultiTHUMOS has a wider range of number of per-class frames and instances (contiguous sequences of a label) annotated than THUMOS. Some action classes like Stand or Run have up to 3.5K instances (up to 18K seconds, or 5.0 hours); others like VolleyballSet or Hug have only 15 and 46 instances (27 and 50 secs) respectively.

action looks visually very different across multiple frames. Finally, MultiTHUMOS has between 27 seconds to 5 hours of annotated video per action class, which challenges action detection algorithms to effectively utilize both small and large amounts of annotated data. In contrast, THUMOS has much less variability, between 3.7 minutes to 1 hour of annotated video per class. Fig. 6 shows the distribution.

4. Technical Approach

Actions in videos exhibit rich patterns, both within a single frame due to action label relations and also across frames as they evolve in time. The desire to elegantly incorporate these cues with state-of-the-art appearance-based models has led to recent works [3, 15, 29] that study combinations of Convolutional Neural Networks (CNN) modeling frame-level spatial appearance and Recurrent Neural Networks (RNN) modeling the temporal dynamics. However, the density of the action labels in our dataset expands the opportunities for more complex modeling at the temporal level. While in principle even a simple instantiation of an ordinary RNN has the capacity to capture arbitrary temporal patterns, it is not necessarily the best model to use in practice. Indeed, our proposed MultiLSTM model extends the recurrent models described in previous work, and our experiments demonstrate its effectiveness.

4.1. LSTM

The specific type of Recurrent architecture that is commonly chosen in previous work is the Long Short-Term Memory (LSTM), which owing to its appealing functional properties has brought success in a wide range of sequence-based tasks such as speech recognition, machine translation and very recently, video activity classification. Let \mathbf{x} be an input sequence (x_1, \dots, x_T) and \mathbf{y} be an output sequence (y_1, \dots, y_T) . An LSTM then maps \mathbf{x} to \mathbf{y} through a series

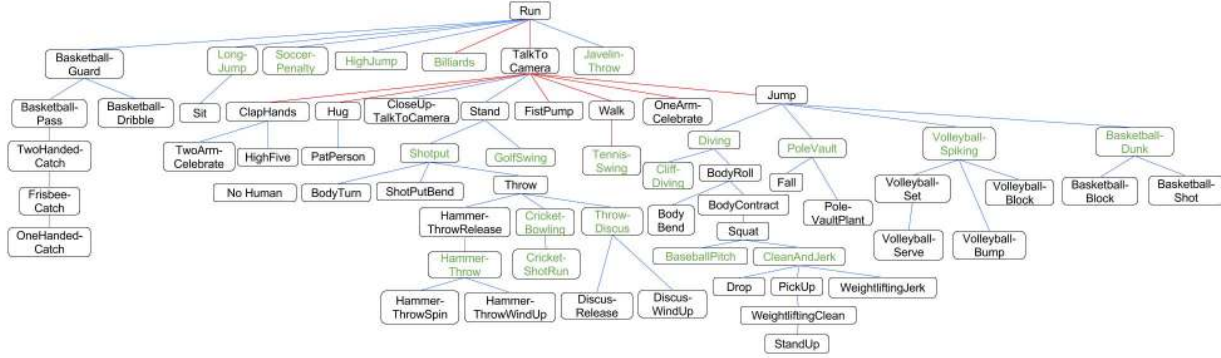


Figure 4: We use the method of [2] to learn the relationships between the 65 MultiTHUMOS classes based on per-frame annotations. Blue (red) means positive (negative) correlation. The 20 original THUMOS classes are in green.

of intermediate representations:

$$i_t = \sigma(W_{xi}x_t + W_{hi}h_{t-1} + b_i) \quad (1)$$

$$f_t = \sigma(W_{xf}x_t + W_{hf}h_{t-1} + b_f) \quad (2)$$

$$o_t = \sigma(W_{xo}x_t + W_{ho}h_{t-1} + b_o) \quad (3)$$

$$g_t = \tanh(W_{xc}x_t + W_{hc}h_{t-1} + b_c) \quad (4)$$

$$c_t = f_t c_{t-1} + i_t g_t \quad (5)$$

$$h_t = o_t \tanh(c_t) \quad (6)$$

$$y_t = W_{hy}h_t + b_y \quad (7)$$

Here c is the “internal memory” of the LSTM, and the gates i , f , o control the degree to which the memory accumulates new input g , attenuates its memory, or influences the hidden layer output h , respectively. Intuitively, the LSTM has the capacity to read and write to its internal memory, and hence maintain and process information over time. Compared to standard RNNs, the LSTM networks mitigate the “vanishing gradients” problem because except for the forget gate, the cell memory is influenced only by additive interactions that can communicate the gradient signal over longer time durations. The architecture is parametrized by the learnable weight matrices W and biases b , and we refer the reader to [5, 3] for further details.

However, an inherent flaw of the plain LSTM architecture is that it is forced to make a definite and final prediction at some time step based on what frame it happens to see at that time step, and its previous context vector.

4.2. MultiLSTM

Our core insight is that providing the model with more freedom in both reading its input and writing its output reduces the burden placed on the hidden layer representation. Concretely, the MultiLSTM expands the temporal receptive field of both input and output connections of an LSTM. These allow the model to directly refine its predictions in retrospect after seeing more frames, and additionally provide direct pathways for referencing previously-seen frames

without forcing the model to maintain and communicate this information through its recurrent connections.

Multilabel Loss. In our specific application setting, the input vectors x_t correspond to the 4096-dimensional fc-7 features of the VGG 16-layer Convolutional Network which was first pretrained on ImageNet and then fine-tuned on our dataset on an individual frame level. We interpret the vectors y_t as the unnormalized log probability of each action class. Since each frame of a video can be labeled with multiple classes, instead of using the conventional softmax loss we sum independent logistic regression losses per class:

$$L(\mathbf{y}|\mathbf{x}) = \sum_{t,c} z_{tc} \log(\sigma(y_{tc})) + (1 - z_{tc}) \log(1 - \sigma(y_{tc}))$$

where y_{tc} is the score for class c at time t , and z_{tc} is the binary ground truth label for class c at time t .

Multiple Inputs with Temporal Attention. In a standard LSTM network, all contextual information is summarized in the hidden state vector. Therefore, the network relies on the memory vector to contain all relevant information about past inputs, without any ability to explicitly revisit past inputs. This is particularly challenging in the context of more complex tasks such as dense, multilabel action detection.

To provide the LSTM with a more direct way of accessing recent inputs, we expand the temporal dimension of the input to be a fixed-length window of frames previous to the current time step (Fig. 7(a)). This allows the LSTM to spend its modeling capacity on more complex and longer-term interactions instead of maintaining summary of the recent frames in case it may be useful for the next few frames. Furthermore, we incorporate a soft-attention weighting mechanism that has recently been proposed in the context of machine translation [1].

Concretely, given a video $\mathbf{V} = \{v_1, \dots, v_T\}$, the input x_i to the LSTM at time step i is now no longer the representation of a single frame v_i , but a weighted combination

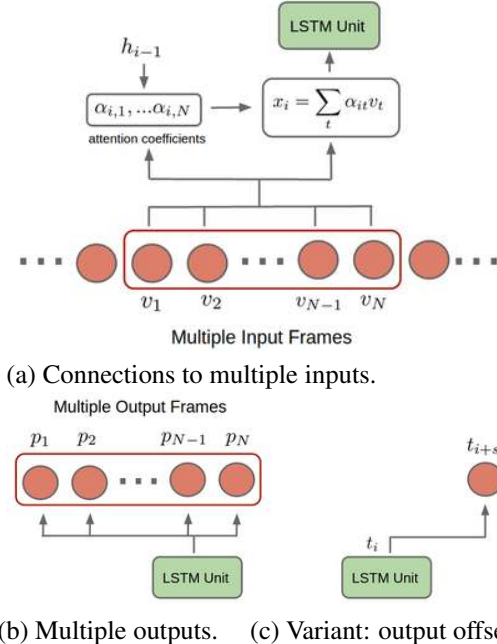


Figure 7: Components of our MultiLSTM model.

$x_i = \sum_t \alpha_{it} v_t$ where t ranges over a fixed-size window of frames previous to i , and α_{it} is the contribution of frame v_t to input x_i as computed by the soft attention model. To compute the attention coefficients α_{it} , we use a model similar to Bahdanau et al. [1]. The precise formulation that worked best in our experiments is:

$$\alpha_{it} \propto \exp(w_{ae}^T [\tanh(W_{ha} h_{i-1}) \odot \tanh(W_{va} v_t)]) \quad (8)$$

Here \odot is element-wise multiplication, $\{w_{ae}, W_{ha}, W_{va}\}$ are learned weights, and α_t is normalized using the softmax function with the interpretation that α_t expresses the relative amount of attention assigned to each frame in the input window. Intuitively, the first term $\tanh(W_{ha} h_{i-1})$ allows the network to look for certain features in the input, while the second term $\tanh(W_{va} v_t)$ allows each input to broadcast the presence/absence of these features. Therefore, the multiplicative interaction followed by the weighted sum with w_{ae} has the effect of quantifying the agreement between what is present in the input and what the network is looking for. Note that the standard LSTM formulation is a special case of this model where all attention is focused on the last input window frame.

Multiple Outputs. Analogous to providing explicit access to a window of frames at the input, we allow the LSTM to contribute to predictions in a window of frames at the output (Fig. 7(b)). Intuitively, this mechanism lets the network refine its predictions in retrospect, after having seen more frames of the input. This feature is related to improvements that can be achieved by use of bi-directional recurrent net-

works. However, unlike bi-directional models our formulation can be used in an online setting where it delivers immediate predictions that become refined with a short time lag.² Given the multiple outputs, we consolidate the predicted labels for all classes c at time t with a weighted average $y_t = \sum_i \beta_{it} p_{it}$ where p_{it} are the predictions at the i th time step for the t th frame, and β_{it} weights the contribution. β_{it} can be learned although in our experiments we use $\frac{1}{N}$ for simplicity to average the predictions. The standard LSTM is a special case, where β is an indicator function at the current time step. In our experiments we use the same temporal windows at the input and output. Similar to the inputs, we experimented with soft attention over the output predictions but did not observe noticeable improvements.

Single Offset Output. We experimented with offset predictions to quantify how informative frames at time t are towards predicting labels at some given offset in time. In these experiments, the network is trained with shifted labels y_{t+s} , where s is a given offset (Fig. 7(c)). In our dense label setting, this type of model additionally enables applications such as action prediction in unconstrained internet video (c.f. [10]). For example, if the input is a frame depicting a person cocking his arm to throw, the model could predict future actions such as Catch or Hit.

5. Experiments

Dataset. We evaluate our MultiLSTM model for dense, multilabel action detection on the MultiTHUMOS dataset. We use the same train and test splits as THUMOS (see Sec. 3 for details) but ignore the background training videos. Clipped training videos (the “Training Data” set in THUMOS) act as weak supervision since they are only labeled with THUMOS classes.

Implementation Details. Our single-frame baseline uses the 16-layer VGG CNN model [27], which achieves near state of the art performance on ILSVRC [23]. The model was pre-trained on ImageNet and fine-tuned on MultiTHUMOS. The input to our LSTM models is the final 4096-dimensional, frame-level fc7 representation.

We use 512 hidden units in the LSTM, and 50 units in the attention component of MultiLSTM. We train the model with an exact forward pass, passing LSTM hidden and cell activations from one mini-batch to the next. However we use approximate backpropagation through time where we only backpropagate errors for the duration of a single mini-batch. Our mini-batches consist of 32 frames (approx. 3.2 seconds), and we use RMSProp [31] to modulate the per-parameter learning rate during optimization.

²A similar behavior can be obtained with a bi-directional model by truncating the hidden state information from future time frames to zero, but this artificially distorts the test-time behavior of the model’s outputs, while our model always operates in the regime it was trained with.




Figure 8: Example timeline of multilabel action detections from our MultiLSTM model compared to a CNN. (best in color)

Performance Measure. We evaluate our models using Average Precision (AP) measured on our frame-level labels. The focus of our work is dense labeling, hence this is the measure we analyze to evaluate the performance of our model. We report AP values for individual action classes as well as mean Average Precision (mAP), the average of these values across the action categories.

To verify that our baseline models are strong, we can obtain discrete detection instances using standard heuristic post-processing.³ Using this post-processing, our single-frame CNN model achieves 32.4 detection mAP with overlap threshold 0.1 on the THUMOS subset of MultiTHUMOS. Since state of the art performance on THUMOS reports 36.6 detection mAP including audio features, this confirms that our single-frame CNN is a reasonable baseline. Hereafter, we compare our models without this post-processing to achieve a comparison of the models’ dense labeling representational ability.

5.1. Action Detection

We first evaluate our models on the challenging task of dense per-frame action labeling on MultiTHUMOS. The MultiLSTM model achieves consistent improvements in mean average precision (mAP) compared to baselines. The basic single-frame CNN achieves 25.4% mAP. A base LSTM network in the spirit of [3] but modified for multilabel action labeling⁴ boosts mAP to 28.1%. Our full MultiLSTM model handily outperforms both baselines with 29.7% mAP. Table 2 additionally demonstrates that each component of our model (input connections, input attention and output connections) is important for accurate action la-

³Concretely, for each class we threshold the frame-level confidences at λ ($\lambda = 0.1$ obtained by cross-validation) to get binary predictions and then accumulate consecutive positive frames into detections. The score of a detection for class C of length L with frame probabilities $p_1 \dots p_L$ is computed as $(\sum_i^L p_i) \times \exp(-\frac{\alpha(L-\mu(C))^2}{\sigma(C)^2})$ where $\mu(C)$ and $\sigma(C)$ are the mean and standard deviation of frame lengths on the training set (for this class) and $\alpha = 0.01$ is obtained by cross-validation.

⁴Specifically, the LSTM is trained using a multilabel loss function and tied hidden context across 32 frame segments, as described in Section 4.2.

Model	THUMOS mAP	MultiTHUMOS mAP
Single-frame CNN [27]	34.7	25.4
Two-stream CNN [26]	36.2	27.6
LSTM	39.3	28.1
LSTM + i	39.5	28.7
LSTM + i + a	39.7	29.1
MultiLSTM	41.3	29.7

Table 2: Per-frame mean Average Precision (mAP) of the MultiLSTM model compared to baselines. Two-stream CNN is computed with single-frame flow. LSTM is implemented in the spirit of [3] (details in Section 4.2). We show the relative contributions of adding first the input connections with averaging (LSTM + i), then the attention (LSTM + i + a) as in Fig. 7(a), and finally the output connections to create our proposed MultiLSTM model (LSTM + i + a + o) as in Fig. 7(b).

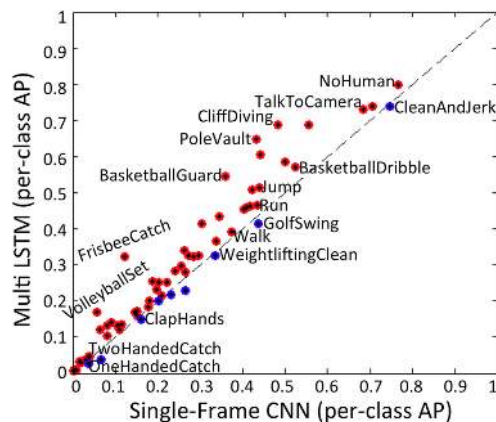


Figure 9: Per-class Average Precision of the MultiLSTM model compared to a single-frame CNN model [27] on MultiTHUMOS. MultiLSTM outperforms the single-frame CNN on 56 out of 65 action classes.

belong. Fig. 9 compares per-class results.

Fig. 8 visualizes some results of MultiLSTM compared to a baseline CNN. For ease of visualization, we binarize



Figure 10: Examples of retrieved sequential actions (correct in green, mistakes in red). Results are shown in pairs: first action frame on the left, second action frame on the right.



Figure 11: Examples of retrieved frames with co-occurring actions (correct in green, mistakes in red).

outputs by thresholding rather than showing the per-frame probabilistic action labels our model produces. The CNN often produces short disjoint detections whereas MultiLSTM effectively makes use of temporal and co-occurrence context to produce more consistent detections.

The multilabel nature of our model and dataset allows us to go beyond simple action labeling and tackle higher-level tasks such as retrieval of video segments containing sequences of actions (Fig. 10) and co-occurring actions (Fig. 11). By learning accurate co-occurrence and temporal relationships, the model is able to retrieve video fragments with detailed action descriptions such as Pass and then Shot or frames with simultaneous actions such as Sit and Talk.



Figure 12: Examples of predicted actions. For each pair of actions, the first one (left) is the label of the current frame and the second one (right) is the predicted label 1 second into the future. Correct predictions are shown in green, and failure cases are shown in red.




Figure 13: Action detection mAP when the MultiLSTM model predicts the action for a past (offset < 0) or future (offset > 0) frame rather than for the current frame (offset $= 0$). The input window of the MultiLSTM model is shown in gray. Thus, the left plot is of a model trained with input from the past, and the right plot is of a model trained with the input window centered around the current frame.

5.2. Action Prediction

Dense multilabel action labeling in unconstrained internet videos is a challenging problem to tackle in and of itself. In this section we go one step further and aim to make predictions about what is likely to happen next or what happened previously in the video. By utilizing the MultiLSTM model with offset (Fig. 7(c)) we are able to use the learned temporal relationships between actions to make inferences about actions likely occurring in past or future frames.

We evaluate the performance of this model as a function of temporal offset magnitude and report results in Fig. 13. The plot on the left quantifies the prediction ability of the model within a 4 second (± 2 second) window, provided an input window of context spanning the previous 1.5 seconds. The model is able to “see the future” – while predicting actions 0.5 seconds in the past is easiest (mAP $\approx 30\%$),

reasonable prediction performance (mAP $\approx 20 - 25\%$) is possible 1-2 seconds into the future. The plot on the right shows the prediction ability of the model using an input context centered around the current frame, instead of spanning only the past. The model is able to provide stronger predictions at past times compared to future times, giving quantitative insight into the contribution of the hidden state vector to providing past context.

Fig. 12 shows qualitative examples of predictions at frames 1 second in the future from the current time. The model is able to correctly infer that a Fall is likely to happen after a Jump, and a BasketballShot soon after a Dribble.

6. Conclusion

In conclusion, this paper presents progress in two aspects of human action understanding. First, we emphasize a broader definition of the task, reasoning about dense, multiple labels per frame of video. We have introduced a new dataset MultiTHUMOS, containing a substantial set of labeled data that we will release to spur research in this direction of action recognition. Second, we develop a novel LSTM-based model incorporating soft attention input-output temporal context for dense action labeling. We show that utilizing this model on our dataset leads to improved accuracy of action labeling and permits detailed understanding of human action.

Acknowledgments

We would like to thank Andrej Karpathy and Amir Zamir for helpful comments and discussion.

References

- [1] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. *CoRR arXiv:1409.0473*, 2014. 5, 6
- [2] M. J. Choi, J. J. Lim, A. Torralba, and A. S. Willsky. Exploiting hierarchical context on a large database of object categories. In *CVPR*, 2010. 4, 5
- [3] J. Donahue, L. A. Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. *CoRR arXiv:1411.4389*, 2014. 2, 3, 4, 5, 7
- [4] G. Gkioxari and J. Malik. Finding action tubes. *CoRR arXiv:1411.6031*, 2014. 3
- [5] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997. 3, 5
- [6] L. Itti et al. A model of saliency-based visual attention for rapid scene analysis. *PAMI*, 20(11):1254–1259, 1998. 3
- [7] Y.-G. Jiang, J. Liu, A. Roshan Zamir, G. Toderici, I. Laptev, M. Shah, and R. Sukthankar. THUMOS challenge: Action recognition with a large number of classes. <http://csrcv.ucf.edu/THUMOS14/>, 2014. 1, 2, 3
- [8] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, 2014. 1, 2
- [9] Y. Ke, R. Sukthankar, and M. Hebert. Event detection in crowded videos. In *ICCV*, 2007. 3
- [10] K. M. Kitani, B. Ziebart, J. D. Bagnell, and M. Hebert. Activity forecasting. In *ECCV*, 2012. 6
- [11] H. Kuehne, A. Arslan, and T. Serre. The language of actions: Recovering the syntax and semantics of goal-directed human activities. In *CVPR*, 2014. 1, 2, 3
- [12] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre. Hmdb: a large video database for human motion recognition. In *ICCV*, 2011. 1, 2
- [13] T. Lan, Y. Wang, and G. Mori. Discriminative figure-centric models for joint action localization and recognition. In *ICCV*, 2011. 3
- [14] E. Mansimov, N. Srivastava, and R. Salakhutdinov. Initialization strategies of spatio-temporal convolutional neural networks. *CoRR arXiv:1503.07274*, 2015. 2
- [15] J. Y.-H. Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, and G. Toderici. Beyond short snippets: Deep networks for video classification. *CoRR arXiv:1503.08909*, 2015. 3, 4
- [16] B. Ni et al. Multiple granularity analysis for fine-grained action detection. In *CVPR*, 2014. 3
- [17] J. C. Niebles, C.-W. Chen, and L. Fei-Fei. Modeling temporal structure of decomposable motion segments for activity classification. In *ECCV*, 2010. 3
- [18] P. Over et al. Trecvid 2011 — an overview of the goals, tasks, data, evaluation mechanisms and metrics. In *Proceedings of TRECVID 2011*, 2011. 2
- [19] H. Pirsiavash and D. Ramanan. Parsing videos of actions with segmental grammars. In *CVPR*, 2014. 3
- [20] R. Poppe. A survey on vision-based human action recognition. *IVC*, 28:976–990, 2010. 2
- [21] M. Rohrbach, S. Amin, M. Andriluka, and B. Schiele. A database for fine grained activity detection of cooking activities. In *CVPR*, 2012. 1, 2, 3
- [22] M. Rohrbach, A. Rohrbach, M. Regneri, S. Amin, M. Andriluka, M. Pinkal, and B. Schiele. Recognizing fine-grained and composite activities using hand-centric features and script data. *CoRR arXiv:1502.06648*, 2015. 3
- [23] O. Russakovsky et al. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 2015. 6
- [24] C. Schuldt, I. Laptev, and B. Caputo. Recognizing human actions: A local svm approach. In *ICPR*, 2004. 2
- [25] N. Shapovalova, M. Raptis, L. Sigal, and G. Mori. Action is in the eye of the beholder: Eye-gaze driven model for spatio-temporal action localization. In *NIPS*, 2013. 3
- [26] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *NIPS*, 2014. 2, 7
- [27] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014. 6, 7

- [28] K. Soomro, A. R. Zamir, and M. Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *CoRR arXiv:1212.0402*, 2012. 1, 2
- [29] N. Srivastava, E. Mansimov, and R. Salakhutdinov. Unsupervised learning of video representations using lstms. *CoRR arXiv:1502.04681*, 2015. 3, 4
- [30] Y. Tian et al. Spatiotemporal deformable part models for action detection. In *CVPR*, 2013. 3
- [31] T. Tieleman and G. E. Hinton. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude., 2012. 6
- [32] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. C3d: Generic features for video analysis. *CoRR arXiv:1412.0767*, 2015. 2
- [33] A. Vahdat, B. Gao, M. Ranjbar, and G. Mori. A discriminative key pose sequence model for recognizing human interactions. In *VS*, 2011. 3
- [34] H. Wang, A. Kläser, C. Schmid, and C.-L. Liu. Action recognition by dense trajectories. In *CVPR*, 2011. 2
- [35] P. Wang et al. Temporal pyramid pooling based convolutional neural networks for action recognition. *CoRR arXiv:1503.01224*, 2015. 2
- [36] D. Weinland, R. Ronfard, and E. Boyer. A survey of vision-based methods for action representation, segmentation and recognition. In *CVIU 115(2)*, pp. 224,241, 2010. 2
- [37] K. Xu et al. Show, attend and tell: Neural image caption generation with visual attention. *CoRR arXiv:1502.03044*, 2015. 3
- [38] J. Yamato, J. Ohya, and K. Ishii. Recognizing human action in time-sequential images using hidden markov model. In *CVPR*, 1992. 3
- [39] L. Yao et al. Video description generation incorporating spatio-temporal features and a soft-attention mechanism. *CoRR arXiv:1502.08029*, 2015. 3
- [40] S. Zha et al. Exploiting image-trained cnn architectures for unconstrained video classification. *CoRR arXiv:1503.04144*, 2015. 2