

Video Event Extraction via Tracking Visual States of Arguments

Guang Yang¹, Manling Li², Jiajie Zhang¹, Xudong Lin³, Heng Ji², Shih-Fu Chang³

¹Tsinghua University

²University of Illinois at Urbana-Champaign

³Columbia University

yangg19@mails.tsinghua.edu.cn, hengji@illinois.edu

Abstract

Video event extraction aims to detect salient events from a video and identify the arguments for each event as well as their semantic roles. Existing methods focus on capturing the overall visual scene of each frame, ignoring fine-grained argument-level information. Inspired by the definition of events as changes of states, we propose a novel framework to detect video events by tracking the changes in the visual states of all involved arguments, which are expected to provide the most informative evidence for the extraction of video events. In order to capture the visual state changes of arguments, we decompose them into changes in pixels within objects, displacements of objects, and interactions among multiple arguments. We further propose Object State Embedding, Object Motion-aware Embedding and Argument Interaction Embedding to encode and track these changes respectively. Experiments on various video event extraction tasks demonstrate significant improvements compared to state-of-the-art models. In particular, on verb classification, we achieve 3.49% absolute gains (19.53% relative gains) in $F_1@5$ on Video Situation Recognition. Our Code is publicly available at <https://github.com/Shinetism/VStates> for research purposes.

Introduction

The ability to comprehend a video requires a structured understanding of events, including what is happening, the objects involved, as well as their semantic roles. Situation Recognition (Yatskar, Zettlemoyer, and Farhadi 2016; Pratt et al. 2020) is a representative task for image event semantic structure understanding, aiming to classify an image to salient events (verbs) and to predict entities (nouns) associated with semantic roles of each event. Video Situation Recognition (Sadhu et al. 2021) extends the capacity of understanding event semantic structures from images to videos, as shown in Figure 1. However, due to the high variation of temporal dynamics and visual features presented in videos, it remains to be a challenging and understudied task.

Traditional image event extraction (Yatskar, Zettlemoyer, and Farhadi 2016; Pratt et al. 2020; Li et al. 2020; Cho, Yoon, and Kwak 2022; Wei et al. 2022) and video event extraction (Medioni et al. 2001; Chen et al. 2021; Sadhu et al.



Figure 1: An example of a video event, with KNOCK as the verb and *man* as TARGET(ARG_1), *bull* as AGENT(ARG_0).

2021) grasps the overall semantics of visual scenes, as a result of the difficulty in locating a specific visual area (i.e., a bounding box) to indicate the happening of an event. This is the case for images; however, additional temporal dynamics of videos provides another perspective to indicate the occurrences of events visually.

Events, in linguistic form and interpretation, are telic, resultative, and indicate changes in states (Dölling, Heyde-Zybatow, and Schäfer 2008). Similarly, we argue that video events can, by nature, be represented by the visual state changes of all involved arguments. We further encode the **argument visual state changes** from three perspectives: (1) Single-frame state is characterized by in-bounding-box *pixel changes* (such as the man’s state changes from a running position to a lying position in Figure 1) and *displacements* of bounding boxes (such as the man moves from left to right in Figure 1). Therefore, we propose **Object State Embedding** (OSE) by encoding both pixel changes and displacements. (2) The motion of objects across multiple frames, which we introduce **Object Motion-aware Embedding** (OME) to aggregate the changes of object states from different frames. (3) The interaction between objects, which we propose to learn an **Multi-Object Interaction Embedding** by encoding the interaction area, which is able to capture the relative position of multiple objects and the pixel changes within intersection background. For example, the bounding boxes of *man* and *bull* are moving towards each other and then being apart, and the background pixels of *grass* also indicate the KNOCK DOWN event.

On top of that, the information of multiple objects and their interactions is aggregated by an **Argument Interaction Encoder**, by *contextualizing* objects with other objects within the same scene, and learning to focus on the objects and interactions that are more indictable. For example, the *man* and the *bull* in Figure 1 are the major objects compared to the *orange ring* in the background. With this module de-

sign, our proposed model learns to rely on the salient objects that are possibly important argument roles for event detection.

As a result of this new paradigm of video event formulation, we classify a salient verb for each video by taking into account fine-grained states of arguments over time, as well as their interactions. The verb classification largely benefits from the joint modeling of multiple arguments, leading to 3.49% absolute gains (19.53% relative gains) in $F_1@5$ according to the Video Situation Recognition validation set. Argument prediction enjoys a further boost due to the argument-level modeling, i.e., 3.97% absolute gains on CIDEr.

Our contributions can be summarized as the follows:

- We propose to formulate video events as the aggregation of visual state changes among multiple arguments. To the best of our knowledge, we are the first work to explore argument-level state changes in video event extraction.
- We propose a new encoding framework by decomposing state changes into in-bounding-box pixel changes, bounding-box displacements, and multi-argument interaction.
- Experiment results on two tasks, verb classification and argument prediction, demonstrate significant improvements compared to state-of-the-art models.

Methodology

Our goal is to identify events (verbs) in a video clip, and to fill in the semantic roles of each predicted verb with entities (nouns). Our main idea is to introduce argument-level information by encoding the state changes of all arguments via Object State Embeddings and Object Motion-aware Embeddings, as well as their interactions through Multi-object Interaction Embeddings (Section). Afterwards, all objects are globally contextualized via an Argument Interaction Encoder (Section) to pay attention to objects that might be of relevance to main arguments for verb classification (Section) and semantic role prediction (Section).

Problem Formulation

Given a video clip containing sampled frames $\{f_i\}_{i=0}^F$ with F as the number of frames, the model is required to output a series of events $\{e\}$ occurring in the video. Each event is represented as a semantic structure

$$e = \{v, \langle r^0, a^0 \rangle, \langle r^1, a^1 \rangle, \dots\}, \quad (1)$$

consisting a verb v and multiple arguments $\{\langle r^k, a^k \rangle\}$ of the event. As defined in the event ontology, each verb v is chosen from a pre-defined verb set \mathcal{V} , and has a pre-defined set of roles $\mathcal{R}(v)$. For example, the role set for the verb KNOCK contains $\mathcal{R}(\text{KNOCK}) = \{\text{AGENT}, \text{TARGET}, \text{SCENE}, \dots\}$. $r^k \in \mathcal{R}(v)$ represents the k^{th} role of the verb v , and a^k denotes the corresponding argument entity, which is described as a word or a phrase without a predefined set of words. For example, the event detected for the video clip in Figure 1 is $v = \text{KNOCK}$, and main arguments are $\langle \text{AGENT}(\text{ARG}_0), \text{gray bull} \rangle$, $\langle \text{TARGET}(\text{ARG}_1), \text{midget in gray hoodie} \rangle$ and $\langle \text{PLACE}(\text{ARG}_{\text{SCENE}}), \text{ground} \rangle$, etc.

Object Visual State Tracking

From the perspective of kinematics, the motion of an object can be decomposed into transitional and rotational movements, which correspond to the body shift and the rotation of the object, respectively. Likewise, we apply this paradigm to videos and decompose the object motion into the *displacement* of the bounding box and the *pixel changes* within that bounding box. As shown in Figure 2, the displacement of the *man* shows that he is moving from left to right and from near to far, while the pixel changes indicate that he is changing from standing to lying down. The motion of the *man* can be well defined by the displacement of its bounding box and the pixel changes. Therefore, we detect objects in the video and propose object state embedding and object motion-aware embedding to encode the displacement and pixel changes of objects.



Figure 2: Bounding boxes of a man in a video.

Object Detection and Tracking We generate object tracklets using the state-of-the-art object tracking model VidVRD (Gao et al. 2021). Taking a raw video as input, it outputs a list of object tracklets $\{o_j\}_{j=0}^O$, where O is the number of recognized objects. Each object tracklet o_j consists a list of its appearances in the video, and each appearance is characterized by a time slot $[t_j^-, t_j^+]$ and the coordinates of the left bottom and the right top corners of the bounding box

$$b_{ji} = \langle (b_{ji}^0, b_{ji}^1), (b_{ji}^2, b_{ji}^3) \rangle \quad (2)$$

at each time stamp inside the time slot. Take Figure 2 as an example. The object *man* is detected on four frames with the bounding boxes extracted from each frame.

Video Encoding Backbone To utilize the superior power of 3D convolutional neural networks and improve computation efficiency, our object state and motion-aware embedding use the state-of-the-art SlowFast model (Feichtenhofer et al. 2019) as the video encoder to get high-level grid features of raw videos. It includes a Slow pathway operating at a low frame rate to encode spatial semantics, and a Fast pathway operating at a high frame rate to capture motion at fine temporal resolution.

The video encoder takes as input slow and fast sampled video frames $f^{\text{slow}} \in \mathbb{R}^{F_1 \times W \times H \times 3}$ and $f^{\text{fast}} \in \mathbb{R}^{F_2 \times W \times H \times 3}$, and outputs slow and fast grid features $g^{\text{slow}} \in \mathbb{R}^{F_1 \times W' \times H' \times d_1}$, $g^{\text{fast}} \in \mathbb{R}^{F_2 \times W' \times H' \times d_2}$:

$$[g^{\text{slow}}, g^{\text{fast}}] = \text{SlowFast}(f^{\text{slow}}, f^{\text{fast}}), \quad (3)$$

where F_1 and F_2 are the number of frames sampled from the video in slow and fast pathways, with the constraint of $F_1 < F_2$. Here, W and H are the width and height of raw video frames, with W' and H' as the width and height of the processed grid features. d_1 and d_2 are the number of channels of slow and fast grid features respectively.




Figure 3: State embedding of the man in the first frame.

Single Frame: Object State Embedding First, in order to encode the states of objects in each frame, we introduce Object State Embedding. As shown in Figure 2, we decompose the object state into in-bounding-box pixels for visual state tracking, as well as the position of the bounding box for displacement tracking.

Given an object o_j in a video frame f_i and its appearance bounding box $b_{ji} = \langle (b_{ji}^1, b_{ji}^2), (b_{ji}^3, b_{ji}^4) \rangle$, we first project the corresponding bounding box to the grid feature space $\hat{b}_{ji} = \langle (\hat{b}_{ji}^1, \hat{b}_{ji}^2), (\hat{b}_{ji}^3, \hat{b}_{ji}^4) \rangle$ by normalizing the coordinates. To encode the state of the object o_j , we mask grid features outside the bounding box and exclusively focus on the object.

The visual state of an object o_j is determined by its current status presented in the frame f_i , such as the running status in Figure 2. Thus, we adopt average pooling over the grid features within the bounding box:

$$p_{ji} = \frac{1}{(\hat{b}_{ji}^2 - \hat{b}_{ji}^0)(\hat{b}_{ji}^3 - \hat{b}_{ji}^1)} \sum_{x=\hat{b}_{ji}^0}^{\hat{b}_{ji}^2-1} \sum_{y=\hat{b}_{ji}^1}^{\hat{b}_{ji}^3-1} g_i[x, y]. \quad (4)$$

where $(\hat{b}_{ji}^0, \hat{b}_{ji}^1)$ and $(\hat{b}_{ji}^2, \hat{b}_{ji}^3)$ represent the left bottom and the right top corners of the bounding box. $g_i[x, y]$ represents the grid feature of the coordinate $\langle x, y \rangle$ in the grid feature space. The grid features g_i are from the slow pathway ($p_{ji}^{\text{slow}} \in \mathbb{R}^{d_1}$) and the fast pathway ($p_{ji}^{\text{fast}} \in \mathbb{R}^{d_2}$).

To track the displacement of the object o_j , we encode its coordinates to track position changes:

$$c_{ji} = W_c \hat{b}_{ji}, \quad (5)$$

where $c_{ji} \in \mathbb{R}^{d_c}$ is the coordinate embedding, with d_c as its dimension. $W_c \in \mathbb{R}^{d_c \times 4}$ is a positional embedding.

As a result, the *Object State Embedding* (OSE) $s_{ji} \in \mathbb{R}^{d+d_c}$ is defined as the concatenation of pooled features p_{ji} and coordinate-based features c_{ji} :

$$s_{ji} = [p_{ji}, c_{ji}].$$

Multiple Frames: Object Motion-aware Embedding To track the state changes of an object o_j across multiple frames $\{f_i\}_{i=0}^F$ along the temporal dimension, we feed object state embeddings of these frames into an *object state aggregator* to capture their differences, as shown in Figure 4. It enables the model to capture the motion of the object o_j when the video unfolds, so we call it *Object Motion-aware Embedding* (OME) m_j :

$$m_j = \text{StateAgg}(\{s_{ji}\}_{i=0}^F), \quad (6)$$




Figure 4: Object Motion-aware Embedding and Object Interaction Embedding.

where $m_j \in \mathbb{R}^{d+d_c}$. The encoding is also done for both slow and fast pathways, i.e., $m_j^{\text{slow}} \in \mathbb{R}^{d_1+d_c}$, $m_j^{\text{fast}} \in \mathbb{R}^{d_2+d_c}$ respectively.

Regarding the object state aggregator, we adopt LSTM (Hochreiter and Schmidhuber 1997) as well as average pooling, and perform further ablation studies on the choice of the aggregation operator.

Multiple Objects: Object Interaction Embedding Apart from the motion of individual objects, we further explore the joint state changes of multiple objects. As one example, the interaction between the *man* and the *bull* in Figure 4 contains the changes of relative positions by moving towards each other, then overlapping, and finally moving far apart. Also, the pixel changes of the interaction area between the *man* and the *bull* provide a context, such as the building, the ground, the falling down circle ring, etc. We therefore propose Object Interaction Embedding to capture and track changes to the interaction area across objects.

To reveal an overall picture of how objects interact in a video frame f_i , we first get the union bounding box \hat{B}_i of these objects:

$$\hat{B}_i = \bigcup_{o_j \in \mathcal{O}_i} \hat{b}_{ji} \quad (7)$$

where \mathcal{O}_i is the set of objects in frame f_i and \cup computes the union of their bounding boxes. Since the interaction area can also be characterized by the pixel changes and position changes, we use similar approach as OME to construct interaction states by the grid features within the interaction area:

$$p'_i = \frac{1}{(\hat{B}_i^2 - \hat{B}_i^0)(\hat{B}_i^3 - \hat{B}_i^1)} \sum_{x=\hat{B}_i^0}^{\hat{B}_i^2-1} \sum_{y=\hat{B}_i^1}^{\hat{B}_i^3-1} g_i[x, y], \quad (8)$$

and the coordinates showing the position of the area:

$$c'_i = W_c \hat{B}_i. \quad (9)$$

Similarly we represent the state of the interaction area of frame f_i via a concatenation operation to gather information of both perspectives:

$$i_i = [p'_i, c'_i], \quad (10)$$

where $i_i \in \mathbb{R}^{d+d_c}$. After that, we employ an interaction aggregator to track the state changes of the interaction area as *Object Interaction Embedding* (OIE) $i \in \mathbb{R}^{d+d_c}$:

$$i = \text{InterAgg}(\{i_i\}_{i=0}^F), \quad (11)$$

where InterAgg follows the same architecture as StateAgg. Notice that we get $i^{\text{slow}} \in \mathbb{R}^{d_1+d_c}$ and $i^{\text{fast}} \in \mathbb{R}^{d_2+d_c}$ through the slow and fast pathways, respectively.

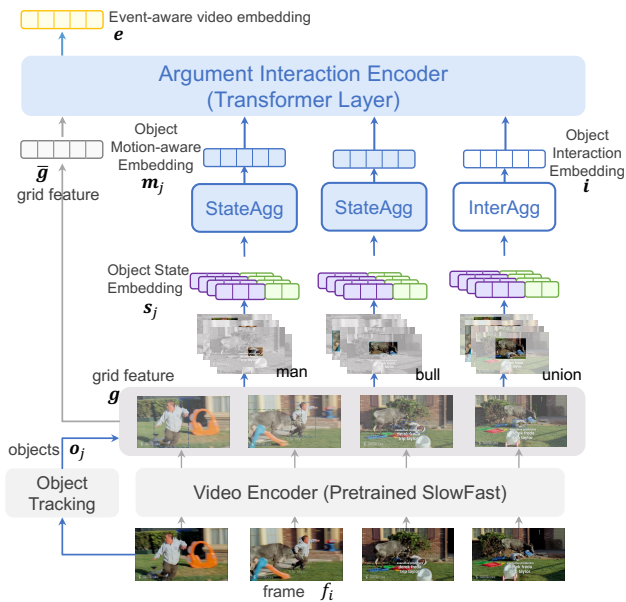


Figure 5: An overview of our event-aware visual embedding.

Argument Interaction Encoder

Considering that not all state changes objects are relevant to the occurrence of the event, we contextualize each object in the video, and ask the model to learn to pay attention to the objects that are most relevant to the video and other objects. Therefore, we employ a Transformer layer (Vaswani et al. 2017) taking as input the overall embedding of entire video, the Object Motion-aware Embedding (OSE) $\{m_j\}_{j=0}^O$ of candidate objects, and the Object Interaction Embedding (OIE) i of the interaction area:

$$e = \text{Trans}(\bar{g}; \{m_j\}_{j=0}^O; i), \quad (12)$$

where $\bar{g} \in \mathbb{R}^d$ is the average-pooled features over all frames, and is used to represent the overall semantics of the video. The maximum number of objects O is set as a hyperparameter, and the objects are ranked in terms of the confidence from the object detector. It enables the output embedding to encode representative information to determine the event semantic structures, so we call it *event-aware visual embedding* e . For each embedding \bar{g} , m_j and i , we concatenate their embeddings from both slow and fast pathways, and further project $e \in \mathbb{R}^{3d_1+3d_2+2d_c}$ to $e \in \mathbb{R}^{d_1+d_2}$ using a linear layer.

Verb Classification

Verb classification is a rudiment motion-related video understanding task to select verbs $\{v\}$ from a predefined event ontology¹. Our event-aware visual embedding e captures the state changes of arguments, which is the most informative feature to identify events. As a result, we classify the video

¹There are 2154 verb-senses defined in the event ontology of the video situation recognition dataset VidSitu (Sadhu et al. 2021).

by feeding its event-aware visual embedding e into a classification layer and train our model with a cross-entropy loss following (Sadhu et al. 2021). The classification layer consists of two linear transformations with a ReLU activation layer in between:

$$v = \text{Softmax}(\mathbf{W}_2 [\text{ReLU}(\mathbf{W}_1 e + \mathbf{b}_1)] + \mathbf{b}_2), \quad (13)$$

which produces a ranked list of the predicted verb v based on the Softmax probabilities. We set the bottleneck dimension to $\frac{1}{2}(d_1 + d_2)$, namely, $\mathbf{W}_1 : \mathbb{R}^{d_1+d_2} \rightarrow \mathbb{R}^{\frac{1}{2}(d_1+d_2)}$ and $\mathbf{W}_2 : \mathbb{R}^{\frac{1}{2}(d_1+d_2)} \rightarrow \mathbb{R}^{|\mathcal{V}|}$, where \mathcal{V} is the set of candidate verbs defined in the event ontology.

Semantic Role Prediction




Figure 6: An overview of our model for semantic role prediction.

The most challenging part of event semantic structure extraction is to recognize the involved objects and their semantic roles. Our event-aware visual embedding e preserves fine-grained argument-level information, so we follow the architecture of state-of-the-art model (Sadhu et al. 2021) but replace its input SlowFast embedding with our event-aware visual embedding e . In detail, the arguments are decoded in a sequence-to-sequence manner shown in Figure 6, where the decoded sequence is the concatenation of the verb and its arguments “ v [Arg₀] a^0 [Arg₁] a^1 ...”, namely, the k^{th} argument a^k is predicted by:

$$a^k = \text{Decoder}(v [\text{Arg}_0] a^0 [\text{Arg}_1] a^1 \cdots [\text{Arg}_k]). \quad (14)$$

Our training objective is cross entropy loss between the generated sequence and the ground-truth sequence. We use teacher forcing strategy to obtain the verb v during training.

Experiments

Dataset

We evaluate our model on VidSitu (Sadhu et al. 2021), the public video dataset providing extensive verb and argument structure annotations for more than 130k video clips. Each video clip is approximately two seconds in length with 10 verbs annotated in the evaluation split, and one verb annotated in the training split. There are 2154 verb-senses defined in the event ontology, and each verb has at least 3 semantic roles. Table 1 shows the dataset statistics. Note that the testing split is hidden and preserved for a leaderboard².

²<https://leaderboard.allenai.org/vidsitu-verbs/submissions/public>

	Train	Valid	Test-Verb	Test-Role
# Clip	118,130	6,630	6,765	7,990
# Verb	118,130	66,300	67,650	79,900
# Role	118,130	19,890	20,295	23,970

Table 1: Data statistics.

Evaluation Metrics

Verb Classification Since each video clip in validation or test split is annotated with 10 verbs, We predict top- K verbs and calculate ranking-based metrics including Accuracy@ K , Recall@ K and $F_1@K$, and set $k = 5$ in experiments following (Sadhu et al. 2021). Also, we report Accuracy@1 as an indicator of its precision in making predictions.

Semantic Role Prediction We follow (Sadhu et al. 2021) to measure the prediction of ARG_0 , ARG_1 , ARG_2 , ARG_{LOC} and ARG_{SCENE} . Since the argument prediction is free-form generation, for every prediction of a semantic role, we compute the CIDEr score (Vedantam, Lawrence Zitnick, and Parikh 2015) by treating the model prediction as candidate sentence and annotations as ground-truth descriptions. Micro-averaged CIDEr score is computed over each individual prediction, while macro-averaged CIDEr score is computed over every verb-sense (CIDEr-Verb) and over argument-types (CIDEr-Arg). Also, we compute ROUGE-L (Lin 2004) for longest common subsequence to show the completeness of generated arguments.

Baselines

State-of-the-art models We compare with the state-of-the-art model in (Sadhu et al. 2021), which contains two variants, including I3D (Carreira and Zisserman 2017a) and SlowFast (Feichtenhofer et al. 2019). For all baselines, we consider the variant with Non-Local blocks (Wang et al. 2018), which has been proved to be more effective according to VidSitu (Sadhu et al. 2021). We provide their performance reported in the paper.

We also include a strong transformer-based video model, TimeSformer (Bertasius, Wang, and Torresani 2021) as an extra baseline to comprehensively assess the performance of our model. As it is not officially evaluated on VidSitu, we tried our best to search for a suitable hyper-parameter setting on this dataset.

For semantic role prediction task, apart from reporting the performance of the state-of-the-art models in (Sadhu et al. 2021), we also include a text-only decoder GPT2 (Radford et al. 2019) as a reference baseline since arguments are predicted via free-form generation.

Unlike verb classification, semantic role prediction results have large variance between difference runs due to free-form generation of argument names. As a result, instead of only reporting performance of a single run, we perform 10 runs and report the average performance (Avg) as well as the standard deviation (Std), as shown in Table 3. For fair com-

parison, we re-run the state-of-the-art model³ using the best-performed SlowFast backbone for 10 times, and report its average and the standard deviation, which is comparable to the performance reported in its paper (Sadhu et al. 2021). We also include the original scores as a reference in Table 3 but dim the color to indicate that they are not fair comparison. Also, we are unable to report the performance on the hidden testing split due to the difficulty of submitting each model 10 times to leaderboard repeatedly, so the average performance is only reported on the validation split.

Ablation models We include variants of our model to explore the effectiveness of our proposed mechanisms, including: (1) **OSE-pixel + OME** by only using Object State Embedding to track *pixel* changes; (2) **OSE-pixel/disp + OME** by using the full setting of Object State Embedding to track both *pixel* changes and *displacements*; (3) **OSE-pixel/disp + OME + OIE** by using the full setting of Object State Embedding and adding Object Interaction Embedding as well.

Note that Object Motion-aware Embedding (OME) is not removable since it is a crucial component to connect multiple frames via a state aggregator, so we add additional experiments and discussions by exploring the different alternatives of the aggregator.

Implementation Details

Verb Classification For verb classification task, We trained our model for 10 epochs and report the model with highest validation $F_1@5$ score. The training took about 20 hour on 4 V100 GPUs, comparable to the original SlowFast baseline. The frame sampling rate and dimensions of grid features are following SlowFast (Feichtenhofer et al. 2019). We keep $d_c = 128$ in our experiments. We set the maximum number of objects as 8, and rank objects in terms of the detection confidence. The learning rate is chosen from $\{10^{-4}, 3 \times 10^{-5}\}$ and the batch size for training is set to 8. We use the Adam optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.99$ and $\epsilon = 10^{-8}$, and no learning rate scheduler is applied to our training. For all pre-trained parameters, we shrink their learning rate by 90%, which keeps the generalizability of the pre-trained encoder.

Semantic Role Prediction The visual event embedding of every video clip is frozen after verb classification training. During the training of semantic role prediction task, we only train the sequence-to-sequence model including the Transformer-based encoder and decoder. The model is trained for 10 epochs and we report the performance with highest validation CIDEr. The learning rate is fixed to 10^{-4} and the batch size is set to 8. We also use the Adam optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.99$ and $\epsilon = 10^{-8}$. Meanwhile, no learning rate scheduler is applied to our training.

Quantitative Analysis

Verb Classification As shown in Table 2, our model outperforms the state-of-the-art models on the leaderboard hidden testing set by 8.19% \sim 8.56% absolute gains on Accuracy@5 and 4.88% \sim 5.05% absolute gains on Recall@5.

³<https://github.com/TheShadow29/VidSitu>

Model	Kinetics	Val				Test			
		Acc@1	Acc@5	Rec@5	F ₁ @5	Acc@1	Acc@5	Rec@5	F ₁ @5
TimeSformer	✓	45.91	79.97	23.61	36.46	-	-	-	-
I3D [†]	✗	30.17	66.83	4.88	9.10	31.43	67.70	5.02	9.35
SlowFast [†]	✗	32.64	69.22	6.11	11.23	33.94	70.54	6.56	12.00
I3D [†]	✓	29.65	60.77	18.21	28.02	29.87	59.10	19.54	29.37
SlowFast [†]	✓	46.79	75.90	23.38	35.75	46.37	75.28	25.78	38.41
Ours (OSE-pixel + OME)	✓	52.75	83.88	28.44	42.48	52.14	83.84	30.66	44.90
Ours (OSE-pixel/disp + OME)	✓	53.32	84.00	28.61	42.68	51.88	83.55	30.83	45.04
Ours (OSE-pixel/disp + OME + OIE)	✓	53.36	83.94	28.72	42.80	52.39	83.47	30.74	44.93

Table 2: Results (%) on verb classification. “Kinetics” denotes whether the model is pre-trained on Kinetics-400 dataset. “OSE-pixel” represents Object State Embedding with in-bounding-box features, and “OSE-pixel/disp” adds displacement capturing. “OME” stands for Object Motion-aware Embedding and “OIE” for Object Interaction Embedding. Results with † is the performance reported in VidSitu paper (Sadhu et al. 2021).

Model	CIDEr		CIDEr-Verb		CIDEr-Arg		ROUGE-L	
	Avg	Std	Avg	Std	Avg	Std	Avg	Std
GPT2 [†]	34.67		42.97		34.45		40.08	
I3D [†]	47.06		51.67		42.76		42.41	
SlowFast [†]	45.52		55.47		42.82		42.66	
SlowFast	44.49 ± 2.30		51.73 ± 2.70		40.93 ± 2.42		40.83 ± 1.27	
Ours (OSE-pixel + OME)	47.82 ± 2.12		54.51 ± 3.00		44.32 ± 2.45		40.91 ± 1.32	
Ours (OSE-pixel/disp + OME)	48.46 ± 1.84		56.04 ± 2.12		44.60 ± 2.33		41.89 ± 1.12	
Ours (OSE-pixel/disp + OME + OIE)	47.16 ± 1.71		53.96 ± 1.32		42.78 ± 2.74		40.86 ± 2.54	

Table 3: Results on semantic role prediction. Results are the average (Avg) over 10 runs with standard deviation (Std) reported. The result with † is the single-run performance reported in VidSitu paper (Sadhu et al. 2021), which is not a fair comparison due to the high variance of free-form generated argument names, so its color is dimmed.

The overall F₁@5 performance achieves 3.32% absolute gains (17.3% relative gains) on testing set and 3.53% absolute gains (19.8% relative gains) on the validation set. It shows the success of using argument state changes to determine events.

Semantic Role Prediction As shown in Table 3, our best performing model typically achieves 4% absolute gains on CIDEr scores, proving the effectiveness of argument-level state tracking in identifying the role of the argument. Additionally, the standard deviation is reduced, demonstrating the robustness of fine-grained argument tracking in event structure parsing.

Pixel Changes vs Displacements Comparison between OSE-pixel and OSE-pixel/disp proves that displacement is important in characterizing object states. It is especially effective in the task of semantic role prediction, showing that tracking the position changes of objects is very important for determining the motion of objects, thus being helpful for detecting semantic roles.

Effect of Object Interaction Embedding (OIE) Verb classification performance is generally improved by capturing interactions between objects, since verbs, by their nature, describe the interactions between objects. However, it does

	Acc@1	Acc@5	Rec@5	F ₁ @5
$O_{\max} = 2$	53.56	83.98	28.95	43.06
$O_{\max} = 4$	53.56	84.10	28.51	42.58
$O_{\max} = 8$	53.35	83.94	28.72	42.80

Table 4: Results (%) of different object numbers on verb classification.

not provide much assistance on argument role prediction, which focuses primarily on each individual argument.

TimeSFormer Comparison Although our model uses a low-performing video encoding backbone SlowFast, the performance is higher on TimeSFormer, demonstrating the effectiveness of our event-aware visual embedding.

Effect of Object Number As shown in Table 5, the number of objects does not significantly affect performance, since the Argument Interaction Encoder calculates the attention across multiple objects and will focus on the objects that may play an important role in event detection.

State Aggregator Operation in OME There is little difference in performance between various alternatives of ob-

	Acc@1	Acc@5	Rec@5	F ₁ @5
LSTM	53.23	83.76	28.57	42.61
Average Pooling	53.32	83.97	28.64	42.71

Table 5: Results (%) of different state aggregator operators on verb classification.

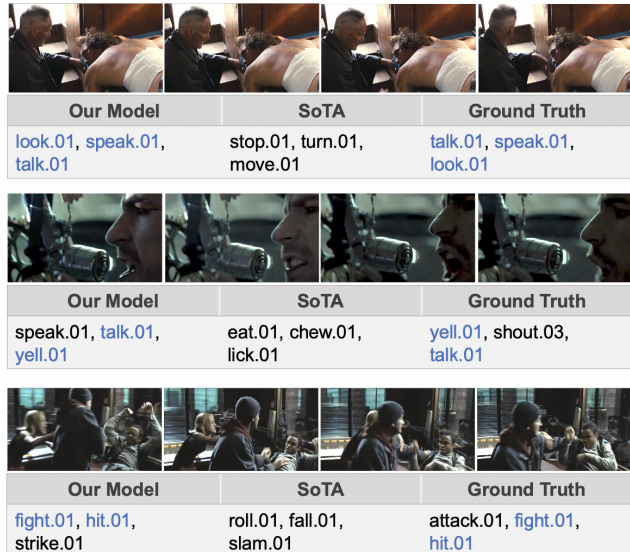


Figure 7: Result comparison between SoTA and our model.

ject state aggregators, and average pooling is slightly more efficient, indicating that a few parameters are sufficient to encode changes of the state of an object.

Qualitative Analysis

Example results in Figure 7 illustrate the successful use of argument state changes as a way of determining events. More examples and analysis are detailed in Appendix due to space limitation.

According to the first example, two objects do not undergo much change in position or pixel, so our model is able to predict events corresponding to LOOK and TALK. However, the SoTA model fails to learn the correlation between object state changes and events, so it predicts events requiring state changes, such as STOP, TURN and MOVE. This example illustrates the reason why tracking changes is important to the detection of events.

The second example illustrates how interactions contribute to the detection of events. A person’s lips are moving as evidenced by pixel changes, so event TALK or CHEW is largely relevant. However, the event CHEW normally requires the interaction between two objects (a person and the food). Since food is not a participant and there is no interaction with food, our model correctly predict the event as TALK. However, SoTA does not take into account fine-grained object-level interaction information, so it fails to predict the correct events but instead prefers CHEW.

The third example demonstrates the important of encoding of pixel changes. It is evident from the position changes of the man on the right side of the video scene that he is moving downward, leading SoTA to prediction of ROLL and FALL. However, his pixel changes indicate that his arm is waving emotionally, which is typically observed in FIGHTING rather than ROLL. As a result, even when displacements match the event type, pixel changes remain extremely useful as an indicator of events.

Related Work

Event Extraction. Extracting events from images/videos (Yatskar, Zettlemoyer, and Farhadi 2016; Pratt et al. 2020; Sadhu et al. 2021), texts (Ji and Grishman 2008; Wang et al. 2019; Liu et al. 2020; Lin et al. 2020b), or multimedia (Li et al. 2020; Chen et al. 2021; Wen et al. 2021; Li et al. 2022) has attracted extensive research efforts. One of the key challenges in event extraction is to model the structural nature (Wang et al. 2019; Li et al. 2020) of events and their associated argument roles. For example, in the text domain, Wang et al. (2019) leverages modular networks to model the relationship between entities of different argument roles. However, to the best of our knowledge, we are the first to explicitly model the states of objects/entities and their relationships in the rich spatial-temporal dynamics of videos.

Video Understanding. Video understanding generally requires the model to capture salient objects (Ikizler-Cinbis and Sclaroff 2010; Gu et al. 2018), their motion (Wang et al. 2016; Carreira and Zisserman 2017a; Shou et al. 2019) and their interactions (Baradel et al. 2018; Gao et al. 2021). End-to-end models (Feichtenhofer et al. 2019; Bertasius, Wang, and Torresani 2021; Lin et al. 2020a) have shown their abilities to capture certain key information for classifying action in videos, where action is usually a single verb or a simple verb phrase (Carreira and Zisserman 2017b). Video event extraction (Sadhu et al. 2021) is a more complicated task than action recognition as it not only requires the model to understand the verb but also the interactions among multiple possible argument roles. Baradel et al. (2018) proposed to learn a reasoning model based on object features for action recognition. However, we first propose to explicitly model the object state, their changes and the interaction of objects for video event extraction.

Conclusions and Future Work

In this work, built upon modern 3D convolutional neural networks, we construct the first event-aware video encoding neural network by tracking arguments visual states, which enables the model to understand event semantic structures. On both verb classification and semantic role prediction tasks in VidSitu, our model achieved a new state of the art, outperforming state-of-the-art 3D convolutional neural networks. The promising future of our event-aware video embedding excites us, as it is applicable to a wide range of event-related video tasks, including action classification, video description, etc. In addition to adapting our model to more tasks, we will also work on tracking argument states in pre-training stages.

Acknowledgements

We thank the anonymous reviewers helpful suggestions.

References

- Baradel, F.; Neverova, N.; Wolf, C.; Mille, J.; and Mori, G. 2018. Object level visual reasoning in videos. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 105–121.
- Bertasius, G.; Wang, H.; and Torresani, L. 2021. Is space-time attention all you need for video understanding? In *ICML*, volume 2, 4.
- Carreira, J.; and Zisserman, A. 2017a. Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, 4724–4733. IEEE Computer Society.
- Carreira, J.; and Zisserman, A. 2017b. Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset. In *CVPR*.
- Chen, B.; Lin, X.; Thomas, C.; Li, M.; Yoshida, S.; Chum, L.; Ji, H.; and Chang, S.-F. 2021. Joint Multimedia Event Extraction from Video and Article. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, 74–88.
- Cho, J.; Yoon, Y.; and Kwak, S. 2022. Collaborative Transformers for Grounded Situation Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19659–19668.
- Dölling, J.; Heyde-Zybatow, T.; and Schäfer, M. 2008. *Event structures in linguistic form and interpretation*, volume 5. Walter de Gruyter.
- Feichtenhofer, C.; Fan, H.; Malik, J.; and He, K. 2019. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, 6202–6211.
- Gao, K.; Chen, L.; Huang, Y.; and Xiao, J. 2021. Video relation detection via tracklet based visual transformer. In *Proceedings of the 29th ACM International Conference on Multimedia*, 4833–4837.
- Gu, C.; Sun, C.; Ross, D. A.; Vondrick, C.; Pantofaru, C.; Li, Y.; Vijayanarasimhan, S.; Toderici, G.; Ricco, S.; Sukthankar, R.; et al. 2018. Ava: A video dataset of spatio-temporally localized atomic visual actions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6047–6056.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation*, 9(8): 1735–1780.
- Ikizler-Cinbis, N.; and Sclaroff, S. 2010. Object, scene and actions: Combining multiple features for human action recognition. In *European conference on computer vision*, 494–507. Springer.
- Ji, H.; and Grishman, R. 2008. Refining Event Extraction through Cross-Document Inference. In *Proceedings of ACL-08: HLT*, 254–262. Columbus, Ohio: Association for Computational Linguistics.
- Li, M.; Xu, R.; Wang, S.; Zhou, L.; Lin, X.; Zhu, C.; Zeng, M.; Ji, H.; and Chang, S.-F. 2022. Clip-event: Connecting text and images with event structures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16420–16429.
- Li, M.; Zareian, A.; Zeng, Q.; Whitehead, S.; Lu, D.; Ji, H.; and Chang, S.-F. 2020. Cross-media Structured Common Space for Multimedia Event Extraction. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2557–2568. Online: Association for Computational Linguistics.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 74–81.
- Lin, X.; Ma, L.; Liu, W.; and Chang, S.-F. 2020a. Context-gated convolution. In *European Conference on Computer Vision*, 701–718. Springer.
- Lin, Y.; Ji, H.; Huang, F.; and Wu, L. 2020b. A Joint End-to-End Neural Model for Information Extraction with Global Features. In *Proc. The 58th Annual Meeting of the Association for Computational Linguistics (ACL2020)*.
- Liu, J.; Chen, Y.; Liu, K.; Bi, W.; and Liu, X. 2020. Event Extraction as Machine Reading Comprehension. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1641–1651. Online: Association for Computational Linguistics.
- Medioni, G.; Cohen, I.; Brémond, F.; Hongeng, S.; and Nevatia, R. 2001. Event detection and analysis from video streams. *IEEE Transactions on pattern analysis and machine intelligence*, 23(8): 873–889.
- Pratt, S.; Yatskar, M.; Weihs, L.; Farhadi, A.; and Kembhavi, A. 2020. Grounded Situation Recognition. In *European Conference on Computer Vision*, 314–332. Springer.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Sadhu, A.; Gupta, T.; Yatskar, M.; Nevatia, R.; and Kembhavi, A. 2021. Visual semantic role labeling for video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5589–5600.
- Shou, Z.; Lin, X.; Kalantidis, Y.; Sevilla-Lara, L.; Rohrbach, M.; Chang, S.-F.; and Yan, Z. 2019. DMC-Net: Generating Discriminative Motion Cues for Fast Compressed Video Action Recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Vedantam, R.; Lawrence Zitnick, C.; and Parikh, D. 2015. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4566–4575.
- Wang, L.; Xiong, Y.; Wang, Z.; Qiao, Y.; Lin, D.; Tang, X.; and Van Gool, L. 2016. Temporal segment networks: Towards good practices for deep action recognition. In *European conference on computer vision*, 20–36. Springer.

- Wang, X.; Girshick, R.; Gupta, A.; and He, K. 2018. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7794–7803.
- Wang, X.; Wang, Z.; Han, X.; Liu, Z.; Li, J.; Li, P.; Sun, M.; Zhou, J.; and Ren, X. 2019. HMEAE: Hierarchical Modular Event Argument Extraction. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 5777–5783. Hong Kong, China: Association for Computational Linguistics.
- Wei, M.; Chen, L.; Ji, W.; Yue, X.; and Chua, T.-S. 2022. Rethinking the Two-Stage Framework for Grounded Situation Recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 2651–2658.
- Wen, H.; Lin, Y.; Lai, T.; Pan, X.; Li, S.; Lin, X.; Zhou, B.; Li, M.; Wang, H.; Zhang, H.; et al. 2021. Resin: A dock-erized schema-guided cross-document cross-lingual cross-media information extraction and event tracking system. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Demonstrations*, 133–143.
- Yatskar, M.; Zettlemoyer, L. S.; and Farhadi, A. 2016. Situation Recognition: Visual Semantic Role Labeling for Image Understanding. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, 5534–5542. IEEE Computer Society.