Look at What I'm Doing: Self-Supervised Spatial Grounding of Narrations in Instructional Videos

Anonymous Author(s)
Affiliation
Address
email

Abstract

We introduce the task of spatially localizing narrated interactions in videos. Key to our approach is the ability to learn to spatially localize interactions with self-supervision on a large corpus of videos with accompanying transcribed narrations. To achieve this goal, we propose a multilayer cross-modal attention network that enables effective optimization of a contrastive loss during training. We introduce a divided strategy that alternates between computing inter- and intra-modal attention across the visual and natural language modalities, which allows effective training via directly contrasting the two modalities' representations. We demonstrate the effectiveness of our approach by self-training on the HowTo100M instructional video dataset and evaluating on a newly collected dataset of localized described interactions in the YouCook2 dataset. We show that our approach outperforms alternative baselines, including shallow co-attention and full cross-modal attention. We also apply our approach to grounding phrases in images with weak supervision on Flickr30K and show that stacking multiple attention layers is effective and, when combined with a word-to-region loss, achieves state of the art on recall-at-one and pointing hand accuracies.

1 Introduction

Content creators often add a voice-over narration to their videos to point out important moments and guide the viewer on where to look. While the timing of the narration in the audio track in the video gives the viewer a rough idea of when the described moment occurs, there is not explicit information on where to look. We, as viewers, naturally infer this information from the narration and often direct our attention to the spatial location of the described moment.

Inspired by this capability, we seek to have a recognition system learn to spatially localize narrated moments without strong supervision. We see a large opportunity for self-supervision as there is an abundance of online narrated video content. In this work, we primarily focus on spatially localizing transcribed narrated interactions in a video, illustrated in Figure 1 (right). Unlike prior phrase grounding work (Figure 1 left), which primarily focuses on matching a noun phrase to an object, our task involves matching entire sentences to regions containing multiple objects and actions. Not only is this task integral to advancing machine perception of our world where information often comes in different modalities, it also has important implications in fundamental vision-and-language research such as robotics, visual question answering, video captioning, and retrieval.

Our task is challenging as we do not know the correspondence between spatial regions in a video and words in the transcribed narration. Moreover, there is a loose temporal alignment of the narration with the described moment and not all words refer to spatial regions (and vice versa). Finally, narrated interactions may be complex with long-range dependencies and multiple actions. In Figure 1 (right), notice how "them" refers to the onions and that there are two actions – "cut" and "add".

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

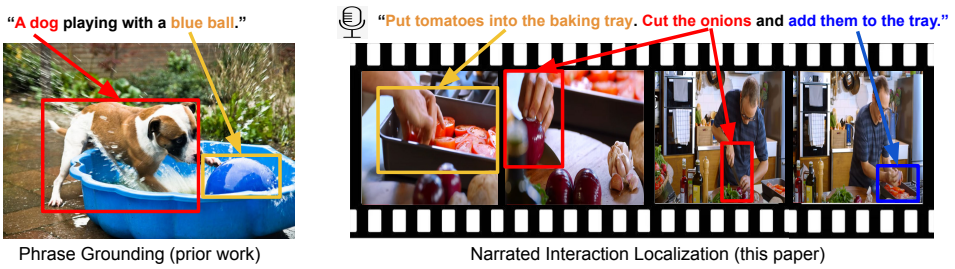

Figure 1: Prior work on *phrase grounding* focuses on localizing single objects in static images (left, example from Flickr30K [29]). In this paper, we propose a new task of spatially *localizing narrated interactions* in video (right). Unlike phrase grounding, interaction grounding may involve multiple objects and actions and use entire sentences rather than nouns or adjective-noun phrases. Our approach self-trains on a large corpus of narrated videos and does not require supervision.

Contrastive learning [27] and computing contextual features via an attention model [37] are natural tools for addressing the aforementioned alignment and long-range cross-modal dependency challenges. While contrastive learning has been successfully applied for temporal alignment of a video clip with a narration [24], the alignment of spatial regions with narrations has not been addressed. Naively applying a joint attention model over the set of features for the two modalities, followed by optimizing a contrastive loss over the two modalities' aggregated contextual features, results in a fundamental difficulty. Recall that an attention computation comprises two main steps: (i) selecting a set of features given a query feature, and (ii) aggregating the selected features via weighted averaging. By selecting and aggregating features over the two modalities during the attention computation, the model has an opportunity to perform *early fusion* and mix text and visual features together. Subsequently optimizing a contrastive loss, the model in theory may select and aggregate features from only one modality to trivially align the contextual representations.

While there are approaches for grounding natural language phrases in images without spatial supervision, the best-performing approaches involve matching individual words to regions [17] or not using any paired text-image supervision at all [39]. This strategy is effective for simpler short phrases, as illustrated in Figure 1 (left). However, consider the interaction "cut the onions and add them to the tray", shown in Figure 1 (right). This interaction is described with more complex, compositional language. Also, as can be seen in the first two frames, there are multiple "onion" objects visible and the correct one depends on the action being applied. We will show that it is difficult to effectively employ a strategy that matches individual words to regions to localize such complex interactions.

To address these challenges, we make the following contributions. First, we learn to localize the spatial location of a narrated interaction from abundantly available narrated instructional videos [25]. Our approach does not require manually collecting sentence descriptions or the locations of the described interactions for training. Here, the automatically transcribed narrations allow for self-supervised learning via alignment with the video clip.

Second, we propose a new approach for directly contrasting aggregated representations from the two modalities while computing joint attention over spatial regions in a video and words in a narration. We show that optimizing a loss over the aggregated, sentence-level representation allows for a global alignment of the described interaction with the video clip, and offers an improvement over optimizing a matching loss over words and regions. To overcome the network learning a trivial solution while directly contrasting jointly attended features from the two modalities during training, we design network attention layers that do not allow feature aggregation (early fusion) across the two modalities. Our strategy involves alternating network layers that compute inter- and intra-modal attention. Our inter-modal attention layer allows features from one modality to select and aggregate features from only the other modality, and not within its own modality. Our intra-modal attention layer selects and aggregates features from within the same modality. This strategy ensures that the output representations do not perform early fusion by aggregating features across the two modalities. In combination with stacking the inter- and intra-modal attention layers, our approach attends jointly and deeply and directly contrasts the resulting contextual representations from the two modalities.

Finally, we introduce an evaluation dataset that provides bounding box annotations for interactions described by natural language sentences. Our dataset is built on the validation split of the YouCook2 dataset [47] and contains approximately 1000 segments of varying durations. We demonstrate our

approach on our collected evaluation dataset of localized interactions and on localizing objects via the YouCook2-BB benchmark where we show that we outperform shallow and full cross-modal attention. We also apply our approach to grounding phrases in images with weak supervision on Flickr30K and show that stacking multiple attention layers is effective with our loss. Furthermore, we show that our loss is complementary with the word-to-region loss of Gupta *et al*. [17], and when combined with it, achieves state of the art on recall-at-one and pointing-hand accuracies.

2 Related Work

Self-supervised learning has been a significant factor in unprecedented progress in natural language processing, resulting in effective and robust learned word representations [26, 13, 42, 28] that achieve state-of-the-art performance on downstream tasks. Recently, it has also garnered a lot of interest in the computer vision community, achieving state-of-the-art performance in unsupervised pre-training of deep visual models. [8, 18, 9] have demonstrated that effective image representations can be learnt simply by contrasting between augmented views of the same image without labels. [44, 32] found that the resulting representations are biased by background pixels and tried to improve foreground object localization via data augmentation or saliency. In video, self-supervised tasks like pace prediction [38] were used for pre-training, while [2] used self-supervision on audio-visual data for tasks like unsupervised speaker localization. In this work, we propose a self-supervised method to spatially localize activity descriptions in video.

Object and action localization Approaches generally leverage a region proposal network (RPN) [36, 16, 15] as well as region-based convolutional neural network (CNNs) to detect objects or localize actions temporally [41] or spatio-temporally [19]. However, they are often trained on curated datasets and are, consequently, limited by a fixed predefined number of object categories in these datasets. [47] collected a dataset of cooking videos labeled with descriptions and temporal segments but did not address spatial grounding of activities. Motivated by this limitation, our proposed approach aims to learn to recognize human interactions with objects that belong to the long-tailed distribution from uncurated and unlabeled online videos. There have been video datasets collected with bounding box annotations with associated natural language descriptions [45]. These datasets provide bounding boxes for noun phrases [11, 45] or for egocentric videos [11]. We collect the first dataset of captioned interactions annotated with spatial bounding boxes for evaluating narrated interactions.

Vision-and-language. Semantic information from natural language has often been exploited to provide an additional source of supervision for learning visual representations. Such approaches have generally leveraged image [20, 22, 29] and video [5] datasets that are annotated with natural language descriptions to learn a joint embedding space, where visual and language representations of semantically-similar pairs are close. Prior approaches have mainly focused on retrieving short clips from a large video corpus [24] or localizing relevant segments within untrimmed videos [5, 7, 10]. However, they do not identify the relevant spatial locations of the described interactions. Some existing works aim to spatially localize a natural language query [31, 43, 21] but focus on objects and train on ground truth bounding boxes, whereas our approach is unsupervised. Object grounding with weak or no supervision is addressed in [10, 4, 6]. Our task is different from object grounding, as it requires not only localizing all objects involved but also identifying when they are undergoing the specified interaction, *c.f.* *"onions"* vs *"add onions to the tray"* in Fig. 1.

Manual annotation of the above-mentioned datasets is usually an expensive and time-consuming process. To alleviate this problem, [24] have proposed to learn robust video representations as well as an effective joint video-text embedding space from a large corpus of unlabeled and uncurated instructional videos [25]. In concurrent work, CLIP [30] trained image-text representations on a dataset of 400 million (image, text) pairs collected from the internet. These efforts to learn from "free" data are similar in spirit to our work, however, we focus on the localization task.

3 Self-Supervised Grounding of Narrated Video Interactions

Given a video clip and a narration, the objective is to learn to localize the relevant spatial regions in the clip which correspond to different interactions described in the narration without bounding box supervision. Formally, we aim to learn language-aligned visual representations for videos such that the distance between features for corresponding narrations and non-matching ones are minimized

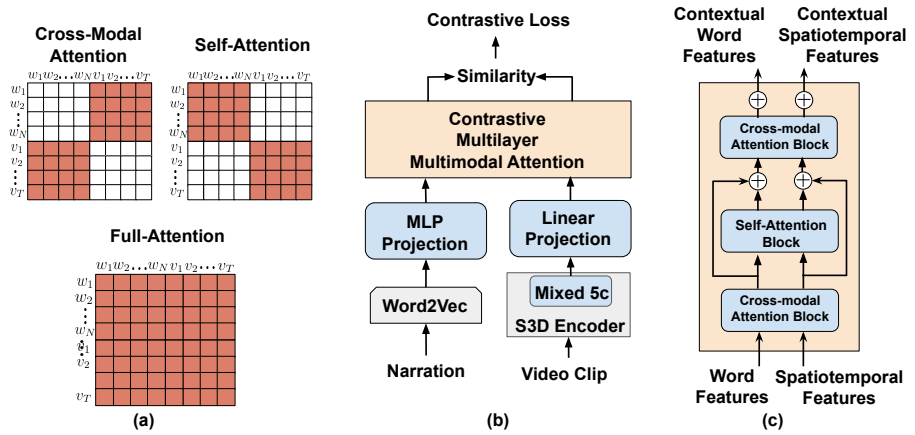

Figure 2: **Approach overview.** (a) Binary masks for computing the different types of cross-modal attention. The shaded regions represent "true" Boolean values where a query feature may select a key. If the mask value is "false", then that corresponding (query, key) pair is not considered in the computation of attention and weighted feature aggregation. (b) Our proposed architecture for computing joint attention over narrations and videos. (c) Our proposed **Contrastive Multilayer Multimodal Attention (CoMMA)** module.

and maximized, respectively. To this end, we propose a novel Contrastive Multilayer Multimodal Attention (CoMMA) module (Figure 2(c)) that allows for direct early interactions between the spatiotemporal regions and words, as well as an effective contrasting of the two modalities' signals.

We define two modalities to have early interactions when they attend to each other to compute their final representations. Conversely, in a late alignment mechanism, the two modalities only interact with each other when the training loss is computed. The intuition underlying our proposed approach is that a deeper cross-modal attention model will encourage a more robust and fine-grained alignment between weakly aligned video clips and sentences. Through repeated early interactions of the two modalities and aggregation of unimodal contextual information, it becomes more discriminative of irrelevant regions or and words. Our approach goes beyond shallow attention architectures [3] and late fusion of the two modalities [24]. Our cross-modal attention module can be easily stacked on top of a base convolutional feature encoder for localization.

Let $c_{x,t}$ be a spatiotemporal feature at spatial location x in frame t and s_i be a word feature for the i-th word in a narration with N words. We obtain features $c_{x,t}$ and s_i via base features that are projected to a common embedding space using linear and MLP layers. We form matrix inputs C_0 and S_0 by stacking the spatiotemporal and word features $c_{x,t}$ and s_i, respectively, as column vectors.

3.1 Contrastive Multilayer Multimodal Attention

The deep attention module comprises alternating layers of a bidirectional cross-modal attention layer followed by a self-attention layer (Figure 2(c)). This architecture choice is based on the reasoning that a hierarchical early interaction process may encourage a model to learn the fine-grained alignment between video regions and narrations while ignoring the irrelevant parts. More specifically, each bidirectional cross-modal layer computes new representations for a target modality with the latent representations from the source modality to learn the relevance of each region to each word or phrase, and vice versa. Additionally, the self-attention layer serves to identify irrelevant words and regions by aggregating contextual information between the augmented unimodal features.

The bidirectional cross-modal attention layer can be formulated as the key/query/value attention mechanism that is commonly used in attention models [37]. The cross-modal setting provides a natural choice for this attention mechanism where the queries are spatiotemporal region features and the keys and values are word features, and vice versa. The attention layer accepts a concatenated sequence of word and region features and a set of binary attention masks. In Figure 2(a), the shaded regions of the masks provide an indication for the attention layer to consider the corresponding elements in the query and key in its computation. Specifically, the cross-modal attention (CA) mask computes the attention over the words with respect to each region and vice versa. Note that this

differs from the full-attention (FA) mask (illustrated in Figure 2(a) bottom) that computes attention over the *entire* multimodal sequence of words and regions. We adopt the attention mechanism in [37] where the attention weights is first computed by measuring the similarity between the query and key before computing a weighted sum of the value representations. Formally, let Q, K, V be the query, key and value matrices, respectively. Let the masked attention mechanism be represented as:

$$Attn(K, Q, V, M) = V \,\mathrm{softmax}\left(\frac{(Q^\top K) \odot M}{\sqrt{D}}\right) \tag{1}$$

where M is any one of the binary masks shown in Figure 2(a), D denotes the dimensionality of the query, and $\odot$ is the Hadamard product.

For simplicity, let matrix $Y_0 = [C_0, S_0]$ denote the stacked input features from both modalities. In addition, let W_K^i, W_Q^i and W_V^i denote the projection matrices for keys, queries and values, respectively, for layer i. We compute the output of the cross-modal attention layer by plugging in the projected video and narration features:

$$Y_{i+1} = CrossAttn(Y_i) = Attn(W_K^i Y_i, W_Q^i Y_i, W_V^i Y_i, M_{CA}) \tag{2}$$

where M_{CA} is the cross-modal attention mask. The similarity scores computed between the queries (regions) and the keys (words), depicted as the normalized dot product between the query and key in Equation (1), act as a soft attention mechanism to measure the relevance of each word with respect to a region. The softmax-normalized scores are multiplied with the value (words) vectors to compute a region-specific narration representation. To illustrate this concept, consider a YouTube video clip about cooking fried chicken. The words 'fried' and 'cook' will have a higher similarity score than 'subscribe' since some regions will be highly relevant to the former. Conversely, the modalities of the keys and queries can be swapped to compute word-specific clip representations.

To aggregate contextual information over the augmented representations of the two modalities, the outputs of the cross-modal attention layer are passed into a self-attention layer. The SA mask is used to ensure that attention is only computed between elements of the same modality. Formally, this computation can be represented as

$$Y_{i+1} = SelfAttn(Y_i) = Attn(W_K^i Y_i, W_Q^i Y_i, W_V^i Y_i, M_{SA}) \tag{3}$$

where M_{SA} is the self-attention mask. Let C_f and S_f denote the outputs of the proposed cross-modal attention module where $C_f \in \mathbb{R}^{DXT}$ is the set of region representations and $S_f \in \mathbb{R}^{DXN}$ is the set of word representations. Mathematically, C_f and S_f are computed as follows:

$$C_f, S_f = f_\theta(C_0, S_0) \tag{4}$$

where f_θ is a function represented by the CoMMA module and parameterized by θ.

Comparison to cross-modal attention models. To evaluate the importance of cross-modal attention modules for localizing narrations, we briefly compare CoMMA against other state-of-the-art variants. [3] is an image-level approach that utilizes features from different levels of an encoder. However, unlike CoMMA, it only utilizes a single round of cross-modal interaction. Another highly relevant work is the Contrastive Bidirectional Transformer (CBT) [35]. Despite the similarities in attention layers, the CBT model computes self-attention over the entire multimodal sequence (similar to the full-attention mask in Figure 2(a)) and aggregates a summary knowledge into a sentinel vector. In contrast, CoMMA allows for repeated early interactions without mixing of features from different modalities. We show in Section 4 that this is critical for localizing from grid features.

3.2 Contrastive loss and inference

We train our proposed cross-modal attention module by contrasting between positive and negative pairs of video clips and narrations. The final clip representation $\hat{C}$ and sentence representation $\hat{S}$ are computed by mean-pooling them over the spatiotemporal regions $\hat{C} = \frac{1}{T}\frac{1}{X}\sum_{i=1}^{T}\sum_{x=1}^{X} C_{fix}$ and words $\hat{S} = \frac{1}{N}\sum_{j=1}^{N} S_{fj}$ respectively. The final contrastive objective is optimized as follows:

$$L_{sentence}(\theta) = -\sum_{i=1}^{n} \log\left(\frac{\exp(\hat{C}^T\hat{S})}{\exp(\hat{C}^T\hat{S}) + \sum\limits_{\hat{S}_{neg}\sim\mathcal{N}_S}\exp(\hat{C}^T\hat{S}_{neg}) + \sum\limits_{\hat{C}_{neg}\sim\mathcal{N}_C}\exp(\hat{C}_{neg}^T\hat{S})}\right) \tag{5}$$

where the negative sets $\mathcal{N}_C$ and $\mathcal{N}_S$ comprises non-corresponding video clip and narration pairs obtained from a given training batch, and n denotes the total number of samples in a training batch. Note that cross-modal attention is applied to negative pairs as well. In contrast to the word-level objective (defined in Section 4) used in [17] which contrasts word features separately, contrasting aggregated features in the sentence loss is more robust to the weak temporal alignment between narrations and videos.

Inference. During inference, given a natural language sentence and a video clip, we aim to localize the salient spatial regions across the frames in the clip. Beginning from the attention weights of the last cross-modal attention layer, we apply Attention Rollout [1] before mean-pooling over all the words to get the final attention heatmap over all spatiotemporal regions. Since the temporal and spatial dimensions of the attention heatmap have been downsampled from the original input resolution, the attention heatmap is temporally and spatially interpolated back to the input resolution. Finally, the mode pixel is used to pinpoint what our model is focusing on given the sentence.

4 Experiments

We evaluate and ablate our method on two video tasks – spatial localization of interactions and objects – and weakly supervised phrase grounding in images.

Word-level loss. We experiment with a word-level loss as used in state-of-the-art phrase grounding [17], which aims to learn an alignment between each word in the narration and all spatial regions. Intuitively, it seeks to maximize the mutual information between each word and the region features if the narration corresponds to the clip. Specifically, the word-level loss is formulated as:

$$L_{word} = -\sum_{i=1}^{n}\sum_{j=1}^{N}\log\left(\frac{\exp\left(S_{fj}^{T}S_{j}^{val}\right)}{\exp\left(S_{fj}^{T}S_{j}^{val}\right)+\sum_{\hat{S}\sim\mathcal{N}_i}\exp\left(\hat{S}_{fj}^{T}S_{j}^{val}\right)}\right) \tag{6}$$

where the negative set $\mathcal{N}_i$ comprises the word features that are attended to by non-corresponding video clips and N denotes the number of words in the narration, and n denotes the total number of samples in a training batch. S_{fj} denotes the representation output of our CoMMA for the j-th word. Finally, for a given word S_j, we compute its value representation as follows: $S_j^{val} = MLP(S_j)$.

4.1 Grounding interactions and objects in video

Self-training dataset. We self-train our proposed model on uncurated instructional videos from the HowTo100M dataset [25]. This dataset contains approximately 1.2 million instructional videos that are sourced from YouTube, encompassing multiple domains such as cooking and hand-crafting. These videos contain narrations that are either uploaded manually by users or the output of an Automatic Speech Recognition (ASR) system. Furthermore, the narrations have been pre-processed to remove all stop words. As noted in prior work [24], there is a weak temporal alignment between the narrations and the video clips, adding another degree of difficulty in training with such videos. The misalignments are due to the nature of instructional videos where an interaction is often narrated before or after demonstrating it. To reduce the computational burden of training over the entire 100M clips, we identify a set of clips from the dataset that roughly aligns with the YouCook2 dataset. To achieve this goal, we extract a list of verb and nouns from the vocabulary of YouCook2 and filter HowTo100M video clips with narrations that contain these words. Our final pre-training set comprises approximately 250,000 video clips. As we will discuss later, note that our base video and language features were trained on the entire HowTo100M dataset, which provides a good starting point for our learning task.

Evaluation datasets. We introduce a new evaluation dataset, YouCook2-Interactions, that provides spatial bounding box annotations for interactions that are described by a natural language sentence. Our dataset is built on the validation split of the YouCook2 dataset [47]. We describe our dataset in full in the supplemental. We evaluate using our collected YouCook2-Interactions dataset for localizing described interactions.

In addition to localizing interactions given a sentence description, we also consider localizing objects given a single object keyword. For localizing single objects, we evaluate on the YouCook2-BB

dataset [46]. The YouCook2-BB dataset is constructed by augmenting the original YouCook2 dataset with bounding box annotations for object locations. In our experiments on YouCook2-BB, we only evaluate our approach on the validation split since the test split annotations are not readily available.

Evaluation criteria. For both video tasks, we evaluate the quality of the output detections using the pointing hand accuracy. Specifically, given a sentence and the ground-truth bounding box, we consider our model to have a "hit" if the pixel with the highest co-attention similarity score lies within the ground-truth bounding box. Otherwise, it is a "miss". The final localization accuracy is computed as follows: Localization Accuracy $= \frac{\text{\# hits}}{\text{\# hits}+\text{\# misses}}$.

Implementation details. We use the Separable 3D CNN-Gated (S3D-G) encoder [40] as well as a shallow language encoder for our video and language models, respectively. To train our proposed model, we set a learning rate of 1e-4 and optimize the model using the AdamW [23] optimizer. An epoch equivalent of steps is used for linear warmup. More details are in the supplemental.

Baselines. In our evaluations, we compare our approach against baselines such as a center prior and optical flow, as well as state-of-the-art image-grounding models. In the center prior, we return the center pixel for each frame. The center prior is a simple yet competitive baseline that performs well due to the nature of these videos, where the main subject is generally in the center of the camera view. In addition, we also implement an optical flow baseline. An optical flow map is generated for each frame using the OpenCV library and the magnitudes are averaged across both spatial directions. Finally, we select the pixel with highest average magnitude as the location prediction.

In addition, we compare against three cross-modal attention baselines. First, we compare against a baseline that leverages full attention with a sentinel vector ("Full attention + sentinel") that predicts the similarity of the two modalities' features, which is then passed to the InfoNCE loss (inspired by CBT [34]). Second, we compare our multilayer model against a shallow cross-modal attention layer ("Shallow attention") inspired by Akbari *et al*. [3]. To maintain fair comparisons, we implement a similar version of their model except we use for the base features the convolutional feature map outputs of the S3D video encoder. Finally, we include comparisons against a state of art phrase grounding model [17]. We apply the model out of the box to our task with the exception of using our features for fair comparison.

Interaction localization. We report in Table 1 results of our approach and the baselines on the interaction localization task. Our proposed approach outperforms the center prior baseline by a significant margin of 20% on the task of localizing interactions given natural language sentences. Surprisingly, computing the region with the maximum degree of motion through optical flow does not provide a good proxy for localization alone. The large performance gain achieved by our proposed approach over the optical flow baseline suggests that tracking regions with maximum movement is insufficient for video understanding. In addition, we observe that our proposed approach outperforms a state-of-the-art phrase grounding model [17] by a large margin. One possible reason is that enforcing an alignment between words and regions is not suitable for grounding interactions in videos, especially when there is a weak temporal alignment between clips and narrations. This is corroborated by a huge drop of 20% in localization accuracy when our CoMMA is trained with the word-level loss instead. As seen in Table 1, our proposed approach outperforms the "full attention + sentinel" baseline by 15%. This result suggests that aggregating contextual information into a sentinel vector causes the model to lose fine-grained information required for localization. We suspect that this is due its cross-modal attention mechanism where features from both modalities are mixed during the interaction process.

Multilayer attention and spatiotemporal self-attention ablations. We report in Table 2 (Left) an ablation over stacking the different attention layers in our network. Notice in rows 1 and 2 that the localization accuracy increases as we stack more attention layers in our network, suggesting that deeper cross-modal attention models help to improve the latent alignment between different modalities. Results from an ablation of the different forms of self-attention over spatiotemporal regions are reported in the right table. It is interesting to observe that by applying self-attention to spatial regions alone, the model performs competitively by obtaining a localization accuracy of 53.88%. We obtain our best model when self-attention is computed over all spatiotemporal regions. This suggests that reasoning about temporal context is important but it has to be done in conjunction with spatial context.

Table 1: **Interaction localization evaluation on our YouCook2-Interactions dataset.** Our approach outperforms baselines, including shallow and full attention.

Approach	Training Loss	Localization Accuracy (%)
Center Prior	-	31.81
Optical flow	-	18.75
Contrastive phrase grounding [17]	Word level	24.04
Shallow Attention	Sentence level	48.30
Full attention + sentinel	Sentence level	39.03
CoMMA (Ours)	Word level	34.91
CoMMA (Ours)	Sentence level	**55.80**

Table 2: **Spatiotemporal self-attention ablation and attention layer ablation.** (Left) We ablate different types of self-attention over spatiotemporal features and find computing full self-attention over spatiotemporal features works best. (Right) We report an ablation of the different attention module components in our model. We find that multiple stacked layers outperform shallower architectures.

Self-Attention Type	Localization Accuracy (%)
Spatial	53.88
Temporal	52.72
Spatial + Temp.	52.15
Spatiotemporal	**55.80**

Cross-Modal Attention 1	Cross-Modal Attention 2	Self-Attention	Localization Accuracy (%)
✓	✗	✗	48.30
✓	✓	✗	51.09
✓	✓	✓	**55.80**

Object localization. Besides localizing full narrated sentences, we also evaluate our model's capability to localize single objects. Specifically, we compare it against the state-of-the-art Not All Frames Are Equal (NAFAE) model [33]. We report results in Table 3. Despite being trained with sentence-level loss, we observe in Table 3 that CoMMA localizes objects corresponding to single noun words. Notably, our proposed approach outperforms NAFAE. We give more details of the NAFAE baseline in the supplemental.

Qualitative results. We show qualitative results in Figure 3. Each frame is overlaid with the computed attention heatmap to depict the visual cues that our model uses to make its prediction. The white circle indicates the location of the center prior and the cross denotes the mode pixel of the computed attention heatmap. Notice how our approach correctly localizes the described interaction.

Table 3: **Object localization evaluation on the YouCook2-BB dataset.** By self-training on narrations, our approach learns to localize objects and outperforms the baseline.

Approach	Full Loc. (%)
NAFAE [33]	46.95
Ours	**59.25**

4.2 Weakly supervised phrase grounding in images

We evaluate our approach on the task of weakly supervised phrase grounding to determine its capability to generalize to localizing noun phrases in images. For this task, we adopt the same setup as in prior work [17, 12, 14] by training on MSCOCO [22] and evaluating on Flickr30k [29]. We report results on the standard recall-at-K (R@K) and pointing accuracy criteria [17]. For training, we optimize the sum of the word-to-region matching loss and our sentence loss: $L(\theta) = L_{word} + \lambda L_{sentence}$. We set λ to be 0.005 in our experiments. We compare our approach against state-of-the-art phrase grounding models. Details of how the model in [17] is extended with our CoMMA architecture as well as a description of the baselines are inclided in the supplemental. We note that the the value projection layer for region features in [17] is omitted to incorporate CoMMA. Finally, we build on the same set of visual and language features from [17] for fair comparison.

Results and analysis. We report results in Table 4. As mentioned above, omitting the region value projection layer results in a performance drop as evidenced by the results obtained by using a single CA layer. Adding a second cross-attention layer generally helps to improve localization accuracy. It leads to notable increases in R@1 and pointing-recall accuracies, surpassing state-of-the-art [17] on those criteria. We note that our full CoMMA model with unimodal self-attention layers hurts performance (R@1 = 48.81%). One possible reason is that optimizing CoMMA on this task is challenging and it may be overfitting. However, it is very promising that adding more layers works out of the box in general on a task that is typically addressed by shallow attention models. In contrast,

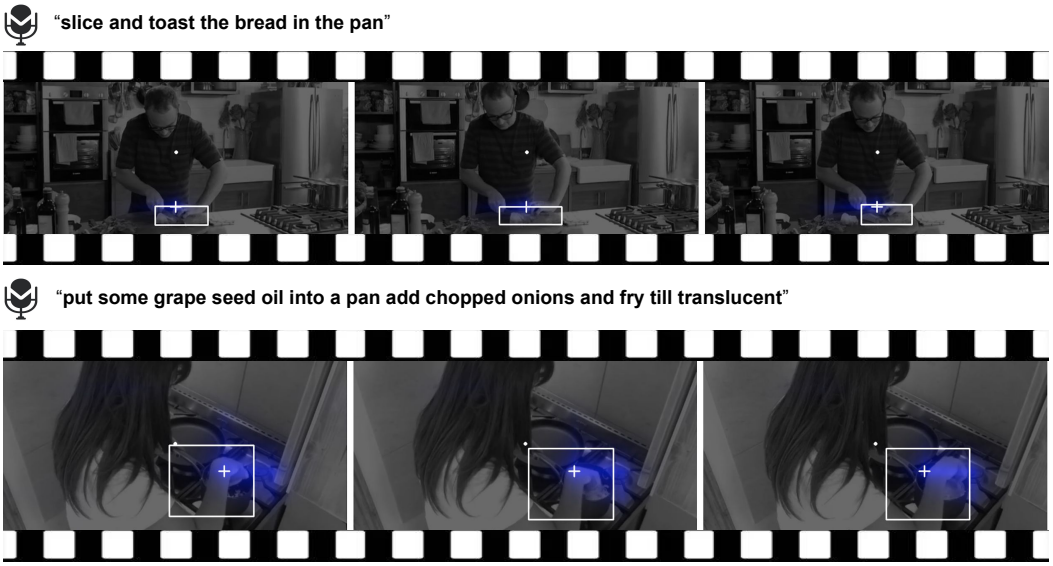

Figure 3: Examples of correct localization predictions by our model. Our approach localizes narrated interactions involving a person's hands and interacted objects.

Table 4: **Weakly supervised phrase grounding experiments on Flickr30K.** Adding additional cross-attention layers improves performance on the R@1 and pointing recall criteria. We include a description of the visual features and baseline approaches in the supplemental. IN and VG denote ImageNet and Visual Genome respectively.

Approach	Visual Features	R@1	R@5	R@10	Pt Recall
Fang *et al* [14]	VGG-cls (IN)	-	-	-	29.00
Akbari *et al* [3]	VGG-cls (IN)	-	-	-	61.66
Akbari *et al* [3]	PNAS Net (IN)	-	-	-	69.19
Align2Ground [12]	Faster-RCNN (VG)	-	-	-	71.00
Gupta *et al* [17]	Faster-RCNN (VG)	51.67	**77.69**	83.25	76.74
ours (word + sent): 1 CA layer	Faster-RCNN (VG)	49.99	76.72	**83.51**	74.97
ours (word + sent): 2 CA layers	Faster-RCNN (VG)	**53.80**	76.69	82.28	**76.78**

we show in the supplemental that stacking multiple layers in our default CoMMA architecture when trained with the sentence-level loss $L_{sentence}$ only results in improved R@K localization accuracies.

Limitations. One limitation is our method operates over coarse feature maps. Future work can aim to refine the localization predictions to compute precise masks for the interactions. However, this could also be used as a platform for further development for surveillance and tracking technology. With this in mind, we plan to release this code judiciously.

Discussion. This work introduces the problem of localizing narrated interactions in uncurated videos without bounding box supervision or priors. To address this, we propose a novel Contrastive Multi-layer Multimodal Attention module that facilitates repeated early interactions between spatiotemporal regions and words to learn their latent alignment. Additionally, we introduce a new evaluation dataset that provides spatial annotations for narrated instructions in videos. Evaluations on the three separate tasks of interaction and object localization in videos as well as phrase grounding in images demonstrate the ability of our model to ground both words and phrases more effectively than the shallow and full-attention baselines. Our approach opens up the possibility of further exploration in other modalities. Learning to associate relevant spatial regions with natural language sentences may also be beneficial to learning richer video representations, particularly with the recent gravitation towards self-supervised representation learning.

References

[1] Samira Abnar and Willem Zuidema. Quantifying attention flow in transformers. *arXiv preprint arXiv:2005.00928*, 2020.

[2] Triantafyllos Afouras, Andrew Owens, Joon Son Chung, and Andrew Zisserman. Self-supervised learning of audio-visual objects from video. *ECCV*, 2020.

[3] Hassan Akbari, Svebor Karaman, Surabhi Bhargava, Brian Chen, Carl Vondrick, and Shih-Fu Chang. Multi-level multimodal common semantic space for image-phrase grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12476–12486, 2019.

[4] Elad Amrani, Rami Ben-Ari, Inbar Shapira, Tal Hakim, and Alex Bronstein. Self-supervised object detection and retrieval using unlabeled videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 954–955, 2020.

[5] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with natural language. In *Proceedings of the IEEE international conference on computer vision*, pages 5803–5812, 2017.

[6] Gedas Bertasius and Lorenzo Torresani. Cobe: Contextualized object embeddings from narrated instructional video. *arXiv preprint arXiv:2007.07306*, 2020.

[7] Jingyuan Chen, Xinpeng Chen, Lin Ma, Zequn Jie, and Tat-Seng Chua. Temporally grounding natural sentence in video. In *Proceedings of the 2018 conference on empirical methods in natural language processing*, pages 162–171, 2018.

[8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.

[9] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020.

[10] Zhenfang Chen, Lin Ma, Wenhan Luo, and Kwan-Yee K Wong. Weakly-supervised spatio-temporally grounding natural sentence in video. *arXiv preprint arXiv:1906.02549*, 2019.

[11] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. The epic-kitchens dataset: Collection, challenges and baselines. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2020.

[12] Samyak Datta, Karan Sikka, Anirban Roy, Karuna Ahuja, Devi Parikh, and Ajay Divakaran. Align2ground: Weakly supervised phrase grounding guided by image-caption alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2601–2610, 2019.

[13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

[14] Hao Fang, Saurabh Gupta, Forrest Iandola, Rupesh K Srivastava, Li Deng, Piotr Dollár, Jianfeng Gao, Xiaodong He, Margaret Mitchell, John C Platt, et al. From captions to visual concepts and back. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1473–1482, 2015.

[15] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015.

[16] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014.

[17] Tanmay Gupta, Arash Vahdat, Gal Chechik, Xiaodong Yang, Jan Kautz, and Derek Hoiem. Contrastive learning for weakly supervised phrase grounding. *arXiv preprint arXiv:2006.09920*, 2020.

[18] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020.

[19] Vicky Kalogeiton, Philippe Weinzaepfel, Vittorio Ferrari, and Cordelia Schmid. Action tubelet detector for spatio-temporal action localization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4405–4413, 2017.

[20] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017.

[21] Zhenyang Li, Ran Tao, Efstratios Gavves, Cees GM Snoek, and Arnold WM Smeulders. Tracking by natural language specification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6495–6503, 2017.

[22] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.

[23] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.

[24] Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, and Andrew Zisserman. End-to-end learning of visual representations from uncurated instructional videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9879–9889, 2020.

[25] Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. Howto100m: Learning a text-video embedding by watching hundred million narrated video clips. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2630–2640, 2019.
[26] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
[27] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2019.
[28] Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*, 2018.
[29] Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proceedings of the IEEE international conference on computer vision*, pages 2641–2649, 2015.
[30] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision, 2021.
[31] Arka Sadhu, Kan Chen, and Ram Nevatia. Video object grounding using semantic roles in language description. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
[32] Ramprasaath R Selvaraju, Karan Desai, Justin Johnson, and Nikhil Naik. Casting your model: Learning to localize improves self-supervised representations. *arXiv preprint arXiv:2012.04630*, 2020.
[33] Jing Shi, Jia Xu, Boqing Gong, and Chenliang Xu. Not all frames are equal: Weakly-supervised video grounding with contextual similarity and visual clustering losses. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10444–10452, 2019.
[34] Chen Sun, Fabien Baradel, Kevin Murphy, and Cordelia Schmid. Contrastive bidirectional transformer for temporal representation learning. In *arXiv*, June 2019.
[35] Chen Sun, Fabien Baradel, Kevin Murphy, and Cordelia Schmid. Learning video representations using contrastive bidirectional transformer. *arXiv preprint arXiv:1906.05743*, 2019.
[36] Jasper RR Uijlings, Koen EA Van De Sande, Theo Gevers, and Arnold WM Smeulders. Selective search for object recognition. *International journal of computer vision*, 104(2):154–171, 2013.
[37] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2017.
[38] Jiangliu Wang, Jianbo Jiao, and Yun-Hui Liu. Self-supervised video representation learning by pace prediction. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, 2020.
[39] Josiah Wang and Lucia Specia. Phrase localization without paired training examples. In *Proceedings of the IEEE international conference on computer vision*, 2019.
[40] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *European Conference on Computer Vision*, 2018.
[41] Huijuan Xu, Abir Das, and Kate Saenko. R-c3d: Region convolutional 3d network for temporal activity detection. In *Proceedings of the IEEE international conference on computer vision*, pages 5783–5792, 2017.
[42] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. *arXiv preprint arXiv:1906.08237*, 2019.
[43] Zhu Zhang, Zhou Zhao, Yang Zhao, Qi Wang, Huasheng Liu, and Lianli Gao. Where does it exist: Spatio-temporal video grounding for multi-form sentences. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10668–10677, 2020.
[44] Nanxuan Zhao, Zhirong Wu, Rynson WH Lau, and Stephen Lin. Distilling localization for self-supervised representation learning. *arXiv preprint arXiv:2004.06638*, 2020.
[45] Luowei Zhou, Yannis Kalantidis, Xinlei Chen, Jason J Corso, and Marcus Rohrbach. Grounded video description. In *CVPR*, 2019.
[46] Luowei Zhou, Nathan Louis, and Jason J Corso. Weakly-supervised video object grounding from text by loss weighting and object interaction. In *British Machine Vision Conference*, 2018.
[47] Luowei Zhou, Chenliang Xu, and Jason Corso. Towards automatic learning of procedures from web instructional videos. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? Yes
   (b) Did you describe the limitations of your work? Yes. See the discussion section.
   (c) Did you discuss any potential negative societal impacts of your work? Yes. See the discussion section.

(d) Have you read the ethics review guidelines and ensured that your paper conforms to them? Yes

2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? N/A
   (b) Did you include complete proofs of all theoretical results? N/A
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? No. We plan to release it after the submission.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? Yes
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? No
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? Yes. We include this in the supplemental.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? Yes
   (b) Did you mention the license of the assets? Yes
   (c) Did you include any new assets either in the supplemental material or as a URL? Yes
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? Yes, we will discuss this in the supplemental.
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? Yes, we will discuss this in the supplemental.
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? Yes
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? N/A
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? Yes, we will discuss this in the supplemental.