Grounded Video Situation Recognition

Anonymous Author(s)
Affiliation
Address
email

Abstract

Dense video understanding requires answering several questions such as *who is doing what to whom, with what, how, why, and where*. Recently, Video Situation recognition (VidSitu) is framed as a task for structured prediction of multiple events, their relationships, their actions and various verb-role pairs attached to descriptive entities. This task poses several challenges in identifying, disambiguating, and co-referencing entities across multiple verb-role pairs, but also faces some challenges of evaluation. In this work, we propose to add spatio-temporal grounding as an essential component of the structured prediction task and present a novel three stage Transformer model, *VideoWhisperer*, that is empowered to make joint predictions. In stage one, we learn contextualised embeddings for video features in parallel with key objects that appear in the video clips to enable fine-grained spatio-temporal reasoning. The second stage sees verb-role queries attend and pool information from object embeddings, localising *answers* to questions posed about the action. The final stage generates these answers as captions to describe each verb-role pair present in the video. Our model operates on a group of events (clips) simultaneously and predicts verbs, verb-role pairs, their nouns, and their grounding on-the-fly. When evaluated on a grounding-augmented version of the VidSitu dataset, we observe a large improvement in entity captioning accuracy, as well as the ability to localize verb-roles without grounding annotations at training time.

1 Introduction

At the end of *The Dark Knight*, we see a short intense sequence[1] that involves Harvey Dent toss a coin while holding a gun followed by sudden action. Holistic understanding of such a video sequence, especially one that involves multiple people, requires predicting more than the action label (*what* verb). For example, we may wish to answer questions such as *who* performed the action (agent), *why* they are doing it (purpose / goal), *how* are they doing it (manner), *where* are they doing it (location), and even *what happens after* (multi-event understanding). While humans are able to perceive the situation and are good at answering such questions, many works often focus on building tools for doing single tasks, *e.g*. predicting actions [1] or detecting objects [2, 3] or image/video captioning [4, 5]. While very important, we are interested in assessing how some of these advances can be combined for a holistic understanding.

A recent and audacious step towards this goal is the work by Sadhu *et al*. [6]. They propose the Video Situation Recognition task (VidSitu) that is a structured prediction over five short clips consisting of three sub-problems: (i) recognizing the salient actions in the short clips; (ii) predicting roles and their entities that are part of this action; and (iii) modeling simple event relations such as enable or cause. Similar to the predecessor image situation recognition (imSitu [7]), VidSitu is annotated using Semantic Role Labelling (SRL) [8]. A video (say 10s) is divided into multiple small events

[1] Spoiler warning! The first 10 seconds of `https://www.youtube.com/watch?v=I6FfPTg1iic`.

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

(∼2s) and each event is associated with a salient action verb (*e.g. hit*). Each verb has a fixed set of roles, *e.g. agent*, *patient*, *tool*, *location*, *manner*, *etc.*, and each role is annotated with a free form text caption, *e.g. agent: Blonde Woman* as shown in Fig. 1.

Grounded VidSitu. VidSitu poses various challenges: long-tailed distribution of both verbs and text phrases, disambiguating the roles, overcoming semantic role-noun pair sparsity, and co-referencing of entities in the entire video. Moreover, there is ambiguity in text phrases that refer to the same unique entity (*e.g.* "man in white shirt" or "man with brown hair"). A model may fail to understand which attributes are important and may bias towards a specific caption (or pattern like shirt color), given the long-tailed distribution. This is exacerbated when multiple entities (*e.g. agent* and *patient*) have similar attributes and the model predicts the same caption for them (see Fig. 1). To remove biases of the captioning module and gauge the model's ability to identify the role, we propose *Grounded Video Situation Recognition* (GVSR) - an extension of the VidSitu dataset to include spatio-temporal grounding. In addition to predicting the captions for the role-entity pairs, we now expect the structured output to contain spatio-temporal localization, currently posed as a weakly-supervised task.

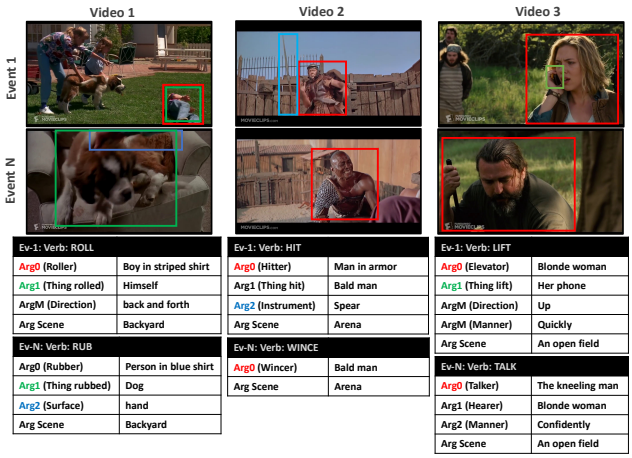

Figure 1: **Overview of GVSR:** Holistic understanding of video requires to localise and recognise all the salient events, and answer questions about it, like who is the agent, who is the patient, what is the location *etc*. GVSR affords this by recognising the actions, their corresponding roles, and localising them in the spatio-temporal domain. This is a challenging task as it requires to disambiguate between several roles that the same entity may take throughout the video, *e.g.* in Video 2 the *bald man* is a *patient* in event 1, but an *agent* in event N. Colored arguments are highlighted in the image with boxes (figure best seen in color).

Joint structured prediction. Previous works [6, 9] modeled the VidSitu tasks separately, *e.g.* the ground-truth verb is fed to the SRL task. This setup does *not* allow for situation recognition on a new video clip without manual intervention. Instead, in this work, we focus on solving three tasks jointly: (i) verb classification; (ii) SRL; and (iii) Grounding for SRL. Note that we drop the original event relation prediction task, as this can be performed later in a decoupled manner similar to [6].

We propose *VideoWhisperer*, a new three-stage transformer architecture that enables video understanding at a global level through self-attention across all video clips, and generates predictions for the three tasks at an event level through localised event-role representations. In the first stage, we use self-attention to align and contextualise 2D object features in addition to event-level video features. These rich features are essential for grounded situation recognition, and are used to predict both the verb-role pairs and entities. In the second stage, a Transformer decoder models the role as a query, and applies cross-attention to find the best elements from the contextualised object features, also enabling visual grounding. Finally, in stage three we generate the captions for each role entity. While the three-stage network disentangles the three tasks, it allows for end-to-end training.

Contributions summary. (i) We present a new framework that combines grounding with SRL for end-to-end Grounded Video Situation Recognition (GVSR). We will release the grounding annotations on the validation set and collaborate with [6] to include them in the evaluation benchmark. (ii) We design a new three-stage transformer architecture for joint verb prediction, semantic-role labelling through caption generation, and weakly-supervised grounding of visual entities. (iii) We propose to use role queries contextualised by video embeddings for SRL, and circumvent the requirement of ground-truth verbs or roles enabling end-to-end GVSR. (iv) At the encoder, we combine object-level features with video-level features, and show multiple advantages in improved weakly-supervised grounding and improved quality of SRL captions leading to a 22 points jump in CIDEr when compared against the video-only baseline [6]. (v) Finally, we present extensive ablation experiments to analyze our model. Our model achieves the state-of-the-art results on the VidSitu benchmark.

2 Related Work

Image Situation Recognition. Situation Recognition in images was first proposed by [10] where they created datasets to understand actions along with localisation of objects and people. Another line of work, imSitu [7] proposed situation recognition via semantic role labelling by leveraging linguistic frameworks, FrameNet [11] and WordNet [12] to formalize situations in the form of verb-role-noun triplets. Recently, grounding has been incorporated with image situation recognition [13] to add a level of understanding for the predicted SRL. Situation recognition requires global understanding of the entire scene, where the verbs, roles and nouns interact with each other to predict a coherent output. Therefore several approaches used CRF [7], LSTMs [13] and Graph neural networks [14] to model the global dependencies among verb and roles. Recently various Transformer [15] based methods have been proposed that claimed huge improvements in performance [16, 17, 18].

Video Situation Recognition. Recently, imSitu was extended to videos as VidSitu [6], a large scale video dataset based on short movie clips spanning multiple events. They used Propbank [8] to annotate the videos with event level verb-role pairs and free form text captions referring to each role. Compared to image situation recognition, VidSRL not only requires understanding the action and the entities involved in a single frame, but also needs to coherently understand the entire video while predicting event-level verb-SRLs and co-referencing the entities participating across events. [6] proposes to use standard video backbones for feature extraction followed by multiple but separate Transformers to model all the tasks individually, using ground-truth of previous the task to model the next. A concurrent work to this submission, [9], proposes to improve upon the video features by pretraining the low-level video backbone using contrastive learning objectives, and pretrain the higher-level video contextualizer using event mask prediction tasks resulting in large performance improvements on SRL. Our goals are different from the above two works, as we propose to learn to perform and predict three tasks simultaneously. We predict verb-role pairs on the fly and design a new role query contextualised by video embeddings to model SRL. This enables end-to-end situation recognition in videos. We also propose to learn contextualised object and video features which enables weakly-supervised grounding for SRL and significantly improves the quality of SRL captions.

Video Understanding. Video understanding is a broad area of research, dominantly involving tasks like action recognition [19, 1, 20, 21, 22, 23], localisation [24, 25], object grounding [26], question answering [27] and, spatio-temporal detection [22]. These tasks involve visual temporal understanding in a sparse uni-dimensional way. In contrast, GVSR involves a hierarchy of tasks, coming together to provide a fixed structure, enabling dense situation recognition. The proposed task requires global video understanding through event level predictions and fine-grained details to recognise all the entities involved, the roles they play, and simultaneously ground them. Note that our work on grounding is different from classical spatio-temporal video grounding [28, 29] or referring expressions based segmentation [30] as they require a text query as input. In our case, both the text and the grounding are predicted jointly by the model.

3 VideoWhisperer for Grounded Video Situation Recognition

We now present the details of our three stage Transformer model, VideoWhisperer. A visual overview is presented in Fig. 2. For brevity, we request the reader to refer to [15] for now popular details of self- and cross-attention layers used in Transformer encoders and decoders.

Preliminaries. Given a video V consisting of several short events $\mathcal{E} = \{e_i\}$, the complete situation in V, is characterized by 3 tasks. (i) Verb classification, requires predicting the action label v_i associated with each event e_i; (ii) Semantic role labelling (SRL), involves guessing the nouns (captions) $\mathcal{C}_i = \{C_{ik}\}$ for various roles $\mathcal{R}_i = \{r | r \in \mathcal{P}(v_i) \forall r \in \mathcal{R}\}$ associated with the verb v_i. $\mathcal{P}$ is a mapping function from verbs to a set of roles based on VidSitu (extended PropBank [8]) and $\mathcal{R}$ is the set of all roles); and (iii) Spatio-temporal Grounding of each visual role-noun prediction C_{ij} is formulated as selecting one among several bounding box proposals $\mathcal{B}$ obtained from sub-sampled keyframes of the video. We evaluate this against ground-truth annotations done at a keyframe level.

3.1 Contextualized Video and Object Features (Stage 1)

GVSR is a challenging task, that requires to coherently model spatio-temporal information to understand the salient action, determine the semantic role-noun pairs involved with the action,

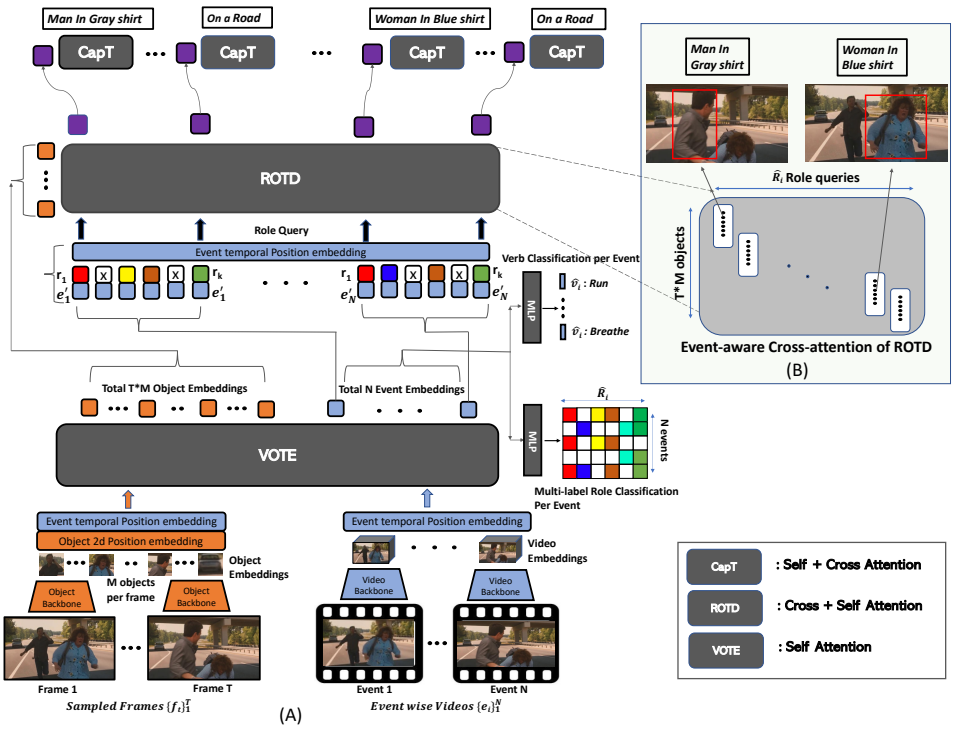

Figure 2: **VideoWhisperer**: We present a new 3-stage Transformer for GVSR. Stage-1 learns the contextualised object and event embeddings through a video-object transformer (VOTE), that is used to predict the verb-role pairs for each event. Stage-2 models all the predicted roles by creating role queries contextualised by event embeddings, and attends to all the object proposals through a role-object transformer (ROTD) to find the best entity that represents a role. The output embeddings are fed to CapT to generate captions for each role. The cross-attention of ROTD ranks all the object proposals enabling localization of each role.

and simultaneously localise them. Different from previous works that operate only on event level video features, we propose to model both the event and object level features simultaneously. We use a pretrained video backbone ϕ_{vid} to extract event level video embeddings $\mathbf{x}_i^e = \phi_{\text{vid}}(e_i)$. For representing objects, we subsample frames $\mathcal{F} = \{f_t\}_{t=1}^T$ from the entire video V. We use a pretrained object detector and extract top M object proposals from every frame. The box locations (along with timestamp) and corresponding features are

$$\mathcal{B} = \{b_{mt}\}, m = [1, \ldots, M], t = [1, \ldots, T], \quad \text{and} \quad \mathbf{x}_{mt}^o = \phi_{\text{obj}}(f_t, b_{mt}) \quad \text{respectively.} \tag{1}$$

The subset of frames associated with an event e_i are computed based on the event's timestamps,

$$\mathcal{F}_i = \{f_t | e_i^{\text{start}} \leq t \leq e_i^{\text{end}}\}. \tag{2}$$

Specifically, at the sub-sampling rate of 1fps, video V of 10s, and events e_i of 2s each, we associate 3 frames with each event (note that border frames are shared). We can extend this mapping to all object proposals based on which frame they appear in and denote this as $\mathcal{B}_i$.

VOTE: Video-Object Transformer Encoder. Since the object and video embeddings come from different spaces, we align and contextualise them with a Transformer encoder [15]. Event-level position embeddings PE_i are added to both representations, event $\mathbf{x}_i^e$ and object $\mathbf{x}_{mt}^o$ ($t \in \mathcal{F}_i$). In addition, 2D object position embeddings PE_{mt} are added to object embeddings $\mathbf{x}_{mt}^o$. Together, they help capture spatio-temporal information. The object and video tokens are passed through multiple self-attention layers to produce contextualised event and object embeddings:

$$[\ldots, \mathbf{o}'_{mt}, \ldots, \mathbf{e}'_i, \ldots] = \text{Transformer}_{\text{VOTE}}([\ldots, \mathbf{x}_{mt}^o + \text{PE}_i + \text{PE}_{mt}, \ldots, \mathbf{x}_i^e + \text{PE}_i, \ldots]). \tag{3}$$

Verb and role classification. Each contextualised event embedding $\mathbf{e}'_i$ is not only empowered to combine information across neighboring events but also focus on key objects that may be relevant. We predict the action label for each event by passing them through a 1-hidden layer MLP,

$$\hat{v}_i = \text{MLP}_e(\mathbf{e}'_i). \tag{4}$$

Each verb is associated with a fixed set of roles based on the mapping $\mathcal{P}(\cdot)$. This prior information is required to model the SRL task. Previous works [6, 9] use ground-truth verbs to model SRL and predict both the roles and their corresponding entities. While this setup allows for task specific modelling, it is not practical in the context of end-to-end video situation recognition. To enable GVSR, we predict the relevant roles for each event circumventing the need for ground-truth verbs and mapped roles. Again, we exploit the contextualised event embeddings and pass them through a role-prediction MLP and perform multi-label role classification. Essentially, we estimate the roles associated with an event as

$$\hat{\mathcal{R}}_i = \{r | \sigma(\text{MLP}_r(\mathbf{e}'_i)) > \theta_{\text{role}}\}\,, \tag{5}$$

where $\sigma(\cdot)$ is the sigmoid function and θ_{role} is a common threshold across all roles (typically set to 0.5). Armed with verb and role predictions, we now look at localizing the role-noun pairs and generating the SRL captions.

3.2 Semantic Role Labelling with Grounding (Stage 2, 3)

A major challenge in SRL is to disambiguate roles as the same object (person) may appear in a different role in the longer video V. For example, if two people are conversing, the *agent* and *patient* roles will switch between *speaker* and *listener* over the course of the video. Another challenge is to generate descriptive and distinctive captions for each role such that they refer to a specific entity. We propose to use learnable role embeddings $\{\mathbf{r}_{ik}\}_{k=1}^{|\mathcal{R}_i|}$ which are capable of learning distinctive role representations. As mentioned earlier, roles such as *agent*, *patient*, *tool*, *location*, *manner*, *etc*. ask further questions about the salient action.

Creating role queries. Depending on the verb each role gets slightly updated. For example, for an action *jump*, the agent would be referred to as the *jumper*. We strengthen the role embeddings by adding the contextualized event embeddings to each role, instead of encoding ground-truth verbs. This eliminates the dependency on the ground-truth verb-role pairs, and enables end to end GVSR. Similar to the first stage (VOTE), we also add event-level temporal positional embeddings to obtain role *query* vectors:

$$\mathbf{q}_{ik} = \mathbf{r}_k + \mathbf{e}'_i + \text{PE}_i\,. \tag{6}$$

Depending on the setting, k can span all roles $\mathcal{R}$, ground-truth roles $\mathcal{R}_i$ or predicted roles $\hat{\mathcal{R}}_i$.

ROTD: Role Object Transformer Decoder. It is hard to achieve rich captions while using features learned for action recognition. Different from prior works [6, 9], we use fine-grained object level representations instead of relying on event-based video features. We now describe the stage two of our VideoWhisperer model, the Transformer decoder for SRL.

Our Transformer decoder uses semantic roles as queries and object proposal representations as keys and values. Through the cross-attention layer, the event-aware role query attends to contextualised object embeddings and finds the best objects that represent each role. We incorporate an event-based attention mask, that limits the roles corresponding to an event to search for objects localised in the same event, while masking out objects from other events. Cross-attention captures local event-level role-object interactions while the self-attention captures the global video level understanding allowing each role to maintain it's own separate identity.

We formulate event-aware cross-attention as follows. We first define the query, key, and value tokens fed to the cross-attention layer as

$$\mathbf{q}'_{ik} = W_Q \mathbf{q}_{ik}, \quad \mathbf{k}'_{mt} = W_K \mathbf{o}'_{mt}, \quad \text{and} \quad \mathbf{v}'_{mt} = W_V \mathbf{o}'_{mt}\,. \tag{7}$$

Here, $W_{[Q|K|V]}$ are learnable linear layers. Next, we apply a mask while computing cross-attention to obtain contextualized role embeddings as

$$\mathbf{r}'_{ik} = \sum_{mt} \alpha_{mt} \mathbf{v}'_{mt}, \quad \text{where} \quad \alpha_{mt} = \text{softmax}_{mt}(\langle \mathbf{q}'_{ik}, \mathbf{k}'_{mt} \rangle \cdot \mathbb{1}(f_t \in \mathcal{F}_i))\,, \tag{8}$$

where $\langle \cdot, \cdot \rangle$ is an inner product and $\mathbb{1}(\cdot)$ is an indicator function with value 1 when true and $-\infty$ otherwise and ensures that the cross-attention is applied only to the boxes $\mathcal{B}_i$, *i.e.* whose frames f_t appear within the same event e_i.

After multiple layers of cross- and self- attention, the role query extracts objects that best represent the entities for each role.

$$[\ldots, \mathbf{z}_{ik}, \ldots] = \text{Transformer}_{\text{ROTD}}([\ldots, \mathbf{q}_{ik}, \ldots; \ldots, \mathbf{o}'_{mt}, \ldots])\,. \tag{9}$$

CapT: Captioning Transformer. The final stage of our model is a caption generation module. Specifically, we use another Transformer decoder [15] whose input context is the output role embedding $\mathbf{z}_{ik}$ from the previous stage and unroll predictions in an autoregressive manner. Different from prior work, captions are generated separately for each role, however, all the necessary information has already been shared through multiple layers of the second stage $\text{Transformer}_{\text{ROTD}}$.

$$\hat{C}_{ik} = \text{Transformer}_{\text{CapT}}(\mathbf{z}_{ik})\,. \tag{10}$$

Grounded Semantic Role Labelling. The entire model is engineered in a way to naturally provide SRL with grounding in a weakly-supervised way, without the ground-truth bounding boxes during training. The cross-attention mechanism of the ROTD Transformer scores and ranks all the objects based on the role-object relevance at every layer. We extract the cross-attention scores α_{mt} for each role k and event e_i from the final layer of $\text{Transformer}_{\text{ROTD}}$, and identify the highest scoring box and the corresponding timestep as

$$\hat{b}^*_m, \hat{b}^*_t = \arg\max_{m,t} \alpha_{mt}\,. \tag{11}$$

3.3 Training and Inference

We briefly discuss how we can train or obtain predictions from our model.

Training. Our entire model can be trained in an end-to-end fashion, with three losses. The first two losses, CrossEntropy and BinaryCrossEntropy, are tapped from the contextualized event embeddings and primarily impact the Video-Object Transformer Encoder

$$L_i^{\text{verb}} = CE(\hat{v}_i, v_i) \quad \text{and} \quad L_i^{\text{role}} = \sum_{r \in \mathcal{R}_i} BCE(r \in \hat{\mathcal{R}}_i, r \in \mathcal{R}_i)\,. \tag{12}$$

The final component is derived from the ground-truth captions and helps produce meaningful SRL outputs. This is also the source of weak supervision for the grounding task,

$$L_{ik}^{\text{caption}} = \sum_w CE(\hat{C}^w_{ik}, C^w_{ik})\,, \tag{13}$$

where the loss is applied in an autoregressive manner to each predicted word w. The combined loss for any training video V is given by

$$\mathcal{L} = \sum_i L_i^{\text{verb}} + \sum_i L_i^{\text{role}} + \sum_{ik} L_{ik}^{\text{caption}}\,. \tag{14}$$

Inference. At test time, we split the video V into similar events e_i and predict verbs $\hat{v}_i$ and roles $\hat{r}_i$ for the same. Here, we have two options: (i) we can use the predicted verb and obtain the corresponding roles using a ground-truth mapping between verbs and roles $\mathcal{P}(\hat{v}_i)$, or (ii) only predict captions for the predicted roles $\hat{\mathcal{R}}_i$. Next, we will showcase the impact of these design choices via several ablations.

4 Experiments

We evaluate our model in two main settings. (i) This setup mimics VidSitu [6], where tasks are evaluated separately. We primarily focus on (a) Verb prediction and (b) Grounded SRL using ground truth verb-role pairs. (ii) End-to-end GVSR, where all the three tasks are modelled together without using ground truth verb-roles.

4.1 Dataset, Metrics, Details

We evaluate our model on the VidSitu [6] dataset that consists of 29k videos (23.6k train, 1.3k val, and others in task-specific test sets) collected from a diverse set of 3k movies. All videos are truncated to 10 seconds, have 5 events spanning 2 seconds each and are tagged with verb and SRL annotations. There are a total of 1560 verb classes and each verb is associated with a fixed set of roles among 11 possible options, however not all are used for evaluation due to noisy annotations (we follow the protocol by [6]). For each role the corresponding value is a free-form caption.

Metrics. For verb prediction, we report Acc@K, *i.e.* event accuracy considering 10 ground-truth verbs and top-K model predictions and Macro-Averaged Verb Recall@K. For SRL we report CIDEr [31], CIDEr-Vb: Macro-averaged across verbs, CIDEr-Arg: Macro averaged across roles, LEA [32], and ROUGE-L [33]. For more details on the metrics pleas refer to [6].

Implementation details. All our implementation is done in Pytorch [34]. We extract event (video) features from a pretrained SlowFast model [1] for video representation (provided by [6]). For object features, we use a FasterRCNN model [35] provided by BUTD [2] pretrained on the Visual Genome dataset [36]. We sample frames at 1 fps from a 10 second video, resulting in $T = 11$ frames. We extract top $M = 15$ boxes from each frame, resulting in 165 objects per video.

All the three Transformers have the same configurations - they have 3 layers with 8 attention heads, and hidden dimension 1024. We use the tokenizer and vocabulary provided by VidSitu [6] which uses byte pair encoding. We have 3 types of learnable embeddings: (i) event position embeddings PE_i with 5 positions corresponding to each event in time; (ii) object localization 2D spatial embedding; and (iii) role embeddings, for each of the 11 roles. The verb classification MLP has a single hidden layer of 2048 d and produces an output across all 1560 verbs. The role classification MLP also has a single hidden layer of 1024 d and produces output in a multi-label setup for all the 11 roles mentioned above. We threshold role prediction scores with $\theta_{\text{role}} = 0.5$.

We use the Adam optimizer [37] with a learning rate of 10^{-4} to train the whole model end-to-end. As we use pretrained features, we train our model on a single RTX-2080 GPU, batch size of 16.

4.2 Grounding SRL: Annotation and Evaluation

As free form captions and their evaluation can be ambiguous, we propose to simultaneously ground each correct role in the spatio-temporal domain. To evaluate grounding performance, we obtain annotations on the validation set. We select the same $T = 11$ frames that are fed to our model sampled at 1fps. For each frame, we ask annotators to see if the visual roles (*agent*, *patient*, *instrument*), can be identified by drawing a bounding box around them on the CVAT tool [38]. For each event i and role k, we consider all valid boxes and create a dictionary of annotations $\mathcal{G}_{ik}$ with keys as frame number and value as bounding box. During prediction, for each role $r \in \hat{\mathcal{R}}_i$, we extract the highest scoring bounding box as in Eq. 11. The Intersection-over-Union (IoU) metric for an event consists of two terms. The first checks if the selected frame appears in the ground-truth dictionary, while the second compares if the predicted box has an overlap greater than θ with the ground-truth annotation,

$$\text{IoU@}\theta = \frac{1}{|\mathcal{R}_i|} \sum_{k=1}^{|\mathcal{R}_i|} \mathbb{1}[\hat{b}_t^* \in \mathcal{G}_{ik}] \cdot \mathbb{1}[\text{IoU}(\hat{b}_m^*, \mathcal{G}_{ik}[t]) > \theta] \,. \tag{15}$$

4.3 Grounded SRL Ablations

We analyze the impact of architecture choices, role query embeddings, and applying a mask in the cross-attention of the role-object decoder (ROTD). All ablations in this section assume access to the ground-truth verb or roles as this allows us to analyze the effect of various design choices. We will report performance over all metrics in the Appendix.

Table 1: Ablations on architecture. All the models use event-aware cross-attention. + indicates stages of the model. RV represents Role-Video decoder, RO is the Role-Object decoder, V is the Video encoder, VO is Video-Object encoder, and C is the Captioning Transformer.

#	Architecture	Query	CIDEr	IoU@0.3	IoU@0.5
1	RV + C	Role + GT-verb Emb	48.03	-	-
2	RO + C	Role + GT-verb Emb	**69.08**	0.21	0.08
3	V + RO + C	Role + Event Emb	68.08	0.23	0.11
4	VO + RO + C	Role + Event Emb	68.23	**0.27**	**0.15**

Architecture Design. We present SRL and grounding results in Table 1. Rows 1 and 2 use a two-stage transformer (ignoring the bottom video-object encoder). As there is no event embedding $\mathbf{e}'_i$, role queries are augmented with ground-truth verb embedding. Using role-object pairs (as is done in ROTD) is critical for good performance on captioning as compared to video representations (CIDEr 69.1 vs. 48.0). Moreover, using objects enables weakly-supervised grounding. We switch from a

two-stage to a three-stage model between rows 1, 2 vs. 3, 4. Rows 2 vs. 4 illustrates the inclusion of the video-object encoder. We see a significant improvement in grounding performance, especially for the harder IoU@0.5 metric (0.08 to 0.15) without affecting captioning performance. Similarly, rows 3 vs. 4 demonstrate the impact of contextualizing object embeddings by events. In particular, using contextualized object representations $\mathbf{o}'_{mt}$ seems to help as compared against base features $\mathbf{x}^o_{mt}$.

Role Query Design. Prior works in situation recognition [6, 16, 17] use verb embeddings to identify entities from both images or videos. In this ablation, we show that instead of learning verb embeddings that only capture the uni-dimensional meaning of a verb and ignore the entities involved, event (or video) embeddings remember details and are suitable for SRL. In fact, Table 2 rows 2 vs. 3 show that event embeddings are comparable and slightly better than GT-verb embeddings when evaluated on Grounding and SRL. Somewhat surprisingly, we see that the role embeddings alone do not perform too badly. We believe this is due to: (i) role embeddings capture the generic meaning of roles like *agent* and *patient* and can generate the correct entities irrespective of the action information; and (ii) The role query attends to object features which are contextualised by video information, so the objects may carry some action information with them.

Table 2: Ablation comparing different role query designs. Architecture: VO + RO + C.

#	Query	CIDEr	IoU@0.3	IoU@0.5
1	Role emb only	66.42	**0.27**	0.11
2	Role + GT-verb Emb	**68.64**	0.25	0.11
3	Role + Event Emb	68.23	**0.27**	**0.15**

Masked Cross-Attention in ROTD. We use event-aware cross-attention to ensure that the roles of an event attend to objects coming from the same event. As seen in Table 3, this reduces the object pool to search from and improves both the SRL and Grounding performance.

Table 3: Impact of masking in ROTD. Architecture: VO + RO + C. Role query: Role + Event emb.

Mask	CIDEr	IoU@0.3	IoU@0.5
No	66.93	0.23	0.09
Yes	**68.23**	**0.27**	**0.15**

4.4 SRL SOTA comparison

In Table 4, we compare our results against VidSitu [6] and a concurrent as-of-yet unpublished work that uses far better features [9]. Nevertheless, we achieve state-of-the-art performance with a 22 points gain in CIDEr score over [6] and a 8 point gain over [9], while using features from [6]. Moreover, our model allows grounding, something not afforded by the previous approaches.

Table 4: Results of SRL with GT verb and role pairs. We reproduce results for VidSitu [6] by teacher-forcing the ground-truth role pairs to make a fair comparison. Results from the concurrent work [9] are as reported in the paper.

Method	CIDEr	C-Vb	C-Arg	R-L	Lea	IoU@0.3	IoU@0.5
SlowFast+TxE+TxD [6]	46.01	56.37	43.58	43.04	**50.89**	-	-
Slow-D+TxE+TxD [9]	60.34	69.12	53.87	43.77	46.77	-	-
VideoWhisperer (Ours)	**68.23**	**74.15**	**61.79**	**45.58**	48.22	**0.27**	**0.15**
Human Level	84.85	91.7	80.15	39.77	70.33	-	-

4.5 GVSR: Towards Joint Prediction

The primary goal of our work is to enable joint prediction of the verb, roles, entities, and grounding.

Verb prediction is an extremely challenging problem due to the long-tail nature of the dataset. In Table 5, we evaluate verb prediction performance when training the model for verb prediction only (rows 1-3) or training it jointly for GVSR (rows 4, 5). Using a simple video only transformer encoder boosts performance over separate modeling of the five event clips (46.8% to 48.8%, rows 1 vs. 2). Including object context through the video-object encoder (row 3) provides an additional boost resulting in highest performance at 49.73% on accuracy@1. A similar improvement is observed in rows 4 vs. 5 (VO vs. V stage one encoder). The reduced performance of rows 4 and 5 as compared against rows 1-3 is primarily because the highest verb accuracy does not coincide with

highest SRL performance. Hence, while the verb accuracy@1 of the GVSR model does reach 49% it degrades subsequently due to overfitting. We observe that the macro-averaged recall@5 is highest for our model, indicating that our model focuses on all verbs rather than just the dominant classes. Addressing this challenge is left for future work.

Understanding role prediction. The verb-role prediction accuracy is crucial for GVSR, since the SRL task is modelled on role-queries. In Table 6 we analyse role prediction in various settings to understand its effect on SRL. Previous work [6] used ground-truth verbs for SRL, while all roles and their entities or values are predicted sequentially. This setting is termed "GT, Pred" (row 2) as it uses the ground-truth verb but predicts the roles. We argue that as the verb-role mappings are a lookup table, this setting is not required. We enable this "GT, GT" setting in [6] by teacher-forcing the GT roles while unrolling role-noun predictions (row 1). Another setting is where the verb is predicted and we obtain roles using the lookup table "Pred, GT map" (row 3). Note that this enables end-to-end SRL, albeit in two steps. The last setting, "Pred, Pred" predicts both verb and role on-the-fly (row 4).

Table 5: Verb prediction performance. Rows 1-3 train only for verb prediction. Rows 4, 5 are trained for GVSR.

#	Architecture	Acc@1	Acc@5	Rec@5
1	Baseline [6]	46.79	75.90	23.38
2	V	48.82	78.01	23.32
3	VO	**49.73**	**78.72**	24.72
4	V + RV + C	40.83	70.73	24.37
5	VO + RV + C	45.06	75.59	**25.25**

Comparing within variants of [6], surprisingly, row 1 does not perform much better than row 2 on CIDEr. We think this may be because the model is trained on GT verbs and is able to predict most of the roles correctly (row 2, Role F1 = 0.88). Subsequently, both rows 3 and 4 show a large performance reduction indicating the over-reliance on ground-truth verb. We see similar trends for our models. Rows 7 and 8 with predicted verb-role pairs lead to reduced SRL performance as compared against rows 5 and 6. Nevertheless, our "Pred, Pred" CIDEr score of 52.3 is still higher than the baseline "GT, GT" at 46.0.

Table 6: Role prediction in various settings. Role F1 is the F1 score averaged over all role classes.

#	Architecture	Verb	Role	V. Acc@1	Role F1	CIDEr
1	VidSitu [6]	GT	GT	-	-	46.01
2		GT	Pred	-	0.88	45.52
3		Pred	GT map	46.79	-	29.93
4		Pred	Pred	46.79	0.47	30.33
5	RO+C	GT	GT	-	-	69.08
6	VO+RO+C	Pred	GT	45.06	-	68.23
7	VO+RO+C	Pred	GT map	45.06	-	51.24
8	VO+RO+C	Pred	Pred	44.05	0.44	52.30

GVSR. Finally, we evaluate our end-to-end model for grounded video situation recognition. In order to enable end-to-end GVSR in [6], we use it in the "Pred, Pred" setting discussed above, that allows verb, role and, SRL predictions. Table 7 shows that our model improves SRL performance over Vidsitu [6] by a margin of 22% on CIDEr score. In addition to that, our model also enables Grounded SRL, that is missing in VidSitu [6]. Due to space restrictions, we will include qualitative results in the Appendix.

Table 7: GVSR: Results for end-to-end situation recognition. Our model architecture is VO + RO + C.

Model	Prediction			Verb	CIDEr	IoU	
	Verb	Role	SRL	Acc@1		0.3	0.5
VidSitu [6]	✓	✓	✓	46.79	30.33	-	-
VideoWhisperer	✓	✓	✓	44.06	52.30	0.13	0.05
	✓	GT	✓	45.06	68.23	0.27	0.15

5 Conclusion

We proposed GVSR as a means for holistic video understanding combining situation recognition - recognizing salient actions, and their semantic role-noun pairs with grounding. We approached this challenging problem by proposing VideoWhisperer, that combines a video-object encoder for contextualised embeddings, video contextualised role query for better representing the roles without the need for ground-truth verbs and an event-aware cross-attention that helps identify the relevant nouns and ranks them to provide grounding. We achieved state-of-the art performance on the VidSitu benchmark with large gains, and also enabled grounding for roles in a weakly-supervised manner.

References

[1] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. SlowFast Networks for Video Recognition. In *International Conference on Computer Vision (ICCV)*, 2019.

[2] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-End Object Detection with Transformers. In *European Conference on Computer Vision (ECCV)*, 2020.

[4] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. ViLBERT: Pretraining Task-Agnostic Visiolinguistic Representations for Vision-and-Language Tasks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

[5] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. VideoBERT: A Joint Model for Video and Language Representation Learning. In *International Conference on Computer Vision (ICCV)*, 2019.

[6] Arka Sadhu, Tanmay Gupta, Mark Yatskar, Ram Nevatia, and Aniruddha Kembhavi. Visual Semantic Role Labeling for Video Understanding. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.

[7] Mark Yatskar, Luke Zettlemoyer, and Ali Farhadi. Situation Recognition: Visual Semantic Role Labeling for Image Understanding. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

[8] Martha Palmer, Daniel Gildea, and Paul Kingsbury. The proposition bank: An annotated corpus of semantic roles. *Computational linguistics*, 31(1):71–106, 2005.

[9] Fanyi Xiao, Kaustav Kundu, Joseph Tighe, and Davide Modolo. Hierarchical Self-supervised Representation Learning for Movie Understanding. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.

[10] Saurabh Gupta and Jitendra Malik. Visual semantic role labeling. *arXiv preprint arXiv:1505.04474*, 2015.

[11] Collin F. Baker, Charles J. Fillmore, and John B. Lowe. The Berkeley FrameNet project. In *Association of Computational Linguistics (ACL)*, 1998.

[12] George A. Miller. Wordnet: A lexical database for english. *Commun. ACM*, 38(11):39–41, nov 1995.

[13] Sarah Pratt, Mark Yatskar, Luca Weihs, Ali Farhadi, and Aniruddha Kembhavi. Grounded Situation Recognition. In *European Conference on Computer Vision (ECCV)*, 2020.

[14] Ruiyu Li, Makarand Tapaswi, Renjie Liao, Jiaya Jia, Raquel Urtasun, and Sanja Fidler. Situation recognition with graph neural networks. In *International Conference on Computer Vision (ICCV)*, 2017.

[15] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

[16] Junhyeong Cho, Youngseok Yoon, Hyeonjun Lee, and Suha Kwak. Grounded situation recognition with transformers. *arXiv preprint arXiv:2111.10135*, 2021.

[17] Meng Wei, Long Chen, Wei Ji, Xiaoyu Yue, and Tat-Seng Chua. Rethinking the two-stage framework for grounded situation recognition. *arXiv preprint arXiv:2112.05375*, 2021.

[18] Junhyeong Cho, Youngseok Yoon, and Suha Kwak. Collaborative transformers for grounded situation recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.

[19] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[20] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *European Conference on Computer Vision (ECCV)*, 2016.

[21] Chen Sun, Abhinav Shrivastava, Carl Vondrick, Kevin Murphy, Rahul Sukthankar, and Cordelia Schmid. Actor-centric relation network. In *European Conference on Computer Vision (ECCV)*, 2018.

[22] Rohit Girdhar, Joao Carreira, Carl Doersch, and Andrew Zisserman. Video action transformer network. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

[23] Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philipp Krahenbuhl, and Ross Girshick. Long-term feature banks for detailed video understanding. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

[24] Tianwei Lin, Xiao Liu, Xin Li, Errui Ding, and Shilei Wen. BMN: Boundary-matching network for temporal action proposal generation. In *International Conference on Computer Vision (ICCV)*, 2019.

[25] Tianwei Lin, Xu Zhao, Haisheng Su, Chongjing Wang, and Ming Yang. BSN: Boundary sensitive network for temporal action proposal generation. In *European Conference on Computer Vision (ECCV)*, 2018.

[26] Arka Sadhu, Kan Chen, and Ram Nevatia. Video object grounding using semantic roles in language description. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.

[27] Zhou Yu, Dejing Xu, Jun Yu, Ting Yu, Zhou Zhao, Yueting Zhuang, and Dacheng Tao. Activitynet-qa: A dataset for understanding complex web videos via question answering. 33:9127–9134, Jul. 2019.

[28] 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 *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.

[29] Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. TubeDETR: Spatio-Temporal Video Grounding with Transformers. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.

[30] Anna Khoreva, Anna Rohrbach, and Bernt Schiele. Video Object Segmentation with Referring Expressions. In *European Conference on Computer Vision (ECCV)*, 2018.

[31] Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.

[32] Nafise Sadat Moosavi and Michael Strube. Which Coreference Evaluation Metric Do You Trust? A Proposal for a Link-based Entity Aware Metric. In *Association of Computational Linguistics (ACL)*, 2016.

[33] Chin-Yew Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, 2004.

[34] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

[35] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2015.

[36] 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.

[37] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[38] CVAT Annotation Tool. `https://github.com/openvinotoolkit/cvat`.

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] **Due to space constraint we will add details in the supplementary material.**
   (c) Did you discuss any potential negative societal impacts of your work? [No]
   (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)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] **In section 4.1, we described all the training details, we will also provide the model code in the supplementary material.**
   (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] Section 4.1, implementation details.
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] **We use the VidSitu [6] dataset and augment it with grounding annotations for the proposed GVSR task.**
   (b) Did you mention the license of the assets? [Yes] **VidSitu [6] is publicly available and licensed under MIT license.**
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] **We will add the grounding annotations over the existing VidSitu dataset in the supplementary material.**
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes] **The VidSitu dataset is publicly available for research purposes. Additional data collection did not require consent.**
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Yes] **The VidSitu data does not contain PII information. It may have content that could be deemed to be offensive due to the nature of movies. Our new annotations do not add to this.**
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] **We will add in the supplementary material.**
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [No] **Annotations were done in-house in the lab. No additional data was collected.**
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [Yes] **We will add it in the supplementary material.**