

# VIEWS: Entity-Aware News Video Captioning

Hammad Ayyubi<sup>1</sup>, Tianqi Liu<sup>2</sup>, Arsha Nagrani<sup>2</sup>, Xudong Lin<sup>1</sup>,  
Mingda Zhang<sup>2</sup>, Anurag Arnab<sup>2</sup>, Feng Han<sup>2</sup>, Yukun Zhu<sup>2</sup>,  
Xuande Feng<sup>1</sup>, Kevin Zhang<sup>1</sup>, Jialu Liu<sup>2</sup>, Shih-Fu Chang<sup>1</sup>,

<sup>1</sup>Columbia University, <sup>2</sup>Google,

Correspondence: hayyubi@cs.columbia.edu

## Abstract

Existing popular video captioning benchmarks and models often produce generic captions for videos that lack specific identification of individuals, locations, or organizations (named-entities). However, in the case of news videos, the setting is more demanding, requiring the inclusion of such named entities for meaningful summarization. Therefore, we introduce the task of directly summarizing news videos into captions that are entities-aware. To facilitate research in this area, we have collected a large-scale dataset named VIEWS (Video NEWS). Within this task, we face challenges inherent to recognizing named entities and navigating diverse, dynamic contexts, all while relying solely on visual cues. To address these challenges, we propose a model-agnostic approach that enriches visual information extracted from videos with context sourced from external knowledge, enabling the generation of entity-aware captions. We validate the effectiveness of our approach across three video captioning models. Additionally, we conduct a critical analysis of our methodology to gain insights into the complexity of the task, the challenges it presents, and potential avenues for future research.

## 1 Introduction

Summarizing video content into a natural language description (i.e. ‘video captioning’) is a crucial aspect of video understanding. It typically requires understanding salient information and dynamics displayed in the video. Existing datasets typically describe videos using largely generic captions: “A woman is giving a speech on stage” (MSR-VTT (Xu et al., 2016)), “An elderly man is playing the piano in front of a crowd” (ActivityNet Captions (Krishna et al., 2017)), “pour in spicy sauce” (YouCook2 (Zhou et al., 2018)) etc. While these captions suffice for daily activity or cooking videos, they are rendered almost meaningless in the domain

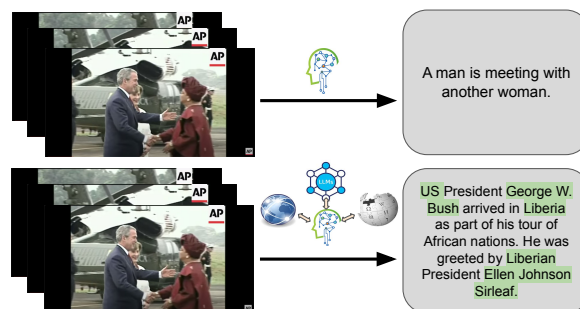


Figure 1: Top: In the traditional Video Captioning task, captions typically do not contain specific named-entities. Bottom: In the proposed Entity-Aware Video Captioning task, the objective is to generate captions that include relevant entities, such as highlighted names and places. This is necessary for effective captioning in the news domain. The task permits utilization of external knowledge sources.

of news videos. As shown in fig. 1, a generic caption saying “a man meets another woman” provides little to no information about: 1) Who is meeting whom? 2) Where? 3) In what context? 4) In relation to which organization? In other words, such captions lack *named entities*.

Indeed, summarizing news videos into entity-aware captions is a challenging task. There has been a line of recent works which additionally utilize the news article paired with the video to generate entity aware captions (Krubiński and Pecina, 2023; Fu et al., 2021; Liu et al., 2021). However, this assumes that a paired news article will always be available for the video. We argue that this is an impractical assumption – there are a number of standalone news video content uploaded online, especially by individuals uninvolved with big news media corporations.

In this paper, we introduce the task of summarizing news videos into entity-aware captions. To support research on this task, we present Video NEWS (VIEWS<sup>1</sup>), comprising a large-scale col-

<sup>1</sup>Data: <https://github.com/hayyubi/views>

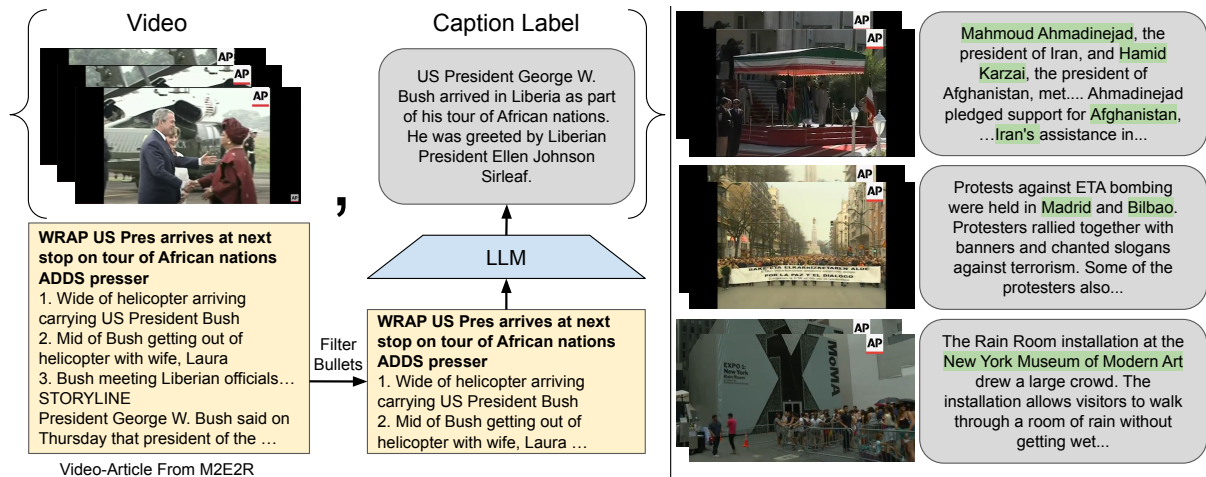


Figure 2: Left: Data Collection Procedure. Right: Samples from our dataset highlighting key properties: captions are rich in entities (colored green), well aligned with the video and diverse (topics shown: politics, conflict & art).

lection of news videos accompanied by entity-rich captions, totaling 144K samples. Unlike previous news video datasets, we rigorously verify and ensure high alignment between the captions and the corresponding videos. Approximately 94% of samples in our training set exhibit high precision, with any remaining precision or alignment issues in the test set manually corrected. Upon acceptance, we will publicly release the dataset.

This task and dataset introduces numerous challenges. Firstly, it requires models to explicitly use visual cues to recognize people, places, events, organizations etc. (named entities) depicted in the video. For instance, in fig. 1, it is required to recognize the people involved in the meeting – George Bush and Ellen Johnson. The landmarks and cultural clues need to be utilized to predict that the meeting is taking place in Liberia. Secondly, understanding the background context from the presented video is a non-trivial task. In fig. 1, this would mean identifying that George Bush is not just touring Liberia, instead, this meeting is part of his African Nations tour. Further, context is incredibly hard to learn – it’s almost unique for every new sample and even changes for the same entities. Lastly, these entities and contexts need to be seamlessly incorporated in the final captioning process.

We propose a novel approach to tackle these challenges. The approach first trains an Entity-Perceiver(EP) to recognize named entities directly from the video. Then, we feed the detected entities to a Knowledge-Extractor (KE) to extract a relevant news data as context from external source. Finally, we leverage recent transformer-based video cap-

Dataset	Domain	#Videos	Entities in Cap.	Cap.-Vid Align ↑
MSR-VTT (Xu et al., 2016)	Open	7.18K	✗	✓
ActyNet Cap (Krishna et al., 2017)	Open	20K	✗	✓
YouCook II (Zhou et al., 2018)	Cooking	2K	✗	✓
TVR/TVC (Lei et al., 2020)	TV show	21.8K	✗*	✓
LSMDC (Rohrbach et al., 2016)	Movie	128K	✗*	✓
MM-AVS (Fu et al., 2021)	News	2.17K	✓	✗
MLASK (Krubiński and Pecina, 2023)	News	41.2K	✓	✗
VIEWS (Ours)	News	144K	✓	✓

Table 1: In comparison to other Video Captioning datasets, VIEWS is the only one to contain entity-rich captions and highly video-aligned captions. Cap: Captions. Vid: Video. s: seconds. ✗\*: grounds entities (Person 1/2/3) but doesn’t name them.

tioning models to integrate the video, the entities, and the context into informative entity-aware captions. We demonstrate the effectiveness of our approach on three state-of-the-art models and study in detail the impact of each of our modules. A detailed analysis leads to many insights and conclusions on where the complexity of the task comes from and how to address them in future works.

## 2 Related Work

**Video Captioning.** Video captioning is a long-standing task that has been explored in various domains: daily activities (Chen and Dolan, 2011; Xu et al., 2016; Krishna et al., 2017), cooking (Zhou et al., 2018; Rohrbach et al., 2015) and movies/TV shows (Bain et al., 2021; Lei et al., 2020; Rohrbach et al., 2016). However, these established video captioning benchmarks do not cover named entities in the real world. This leads to a vacancy in model development as well: recent multimodal Large Language Models (LLMs) (Alayrac et al., 2022) demonstrate impressive performance on traditional

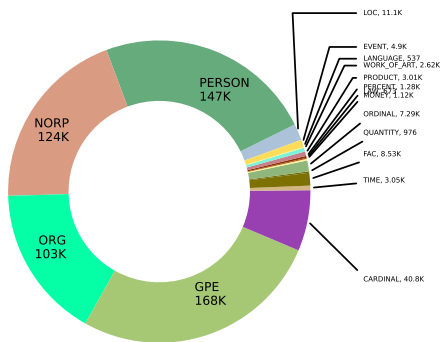


Figure 3: Distribution of entities in VIEWS. NORP: Nationalities or religious or political groups. GPE: geopolitical entities (countries, cities, states). ORG: organizations (companies, agencies, institutions).

benchmarks, but tend to hallucinate named-entities (Ji et al., 2023; Li et al., 2023b). We address this gap through the proposed dataset, VIEWS. VIEWS contains captions that are rich in entities by focusing on the news domain, and encourages entities-aware model development.

**News Content Understanding.** Recent works have attempted to specifically generate entity-aware captions for videos (Fu et al., 2021; Tang et al., 2022; Liang et al., 2023; Whitehead et al., 2018) and images (Biten et al., 2019; Tran et al., 2020; Lu et al., 2018; Zhao et al., 2019). However, they require paired news articles or additional metadata (tags, dates, or named-entities) to do so. This makes these approaches unrealistic in real life, where paired articles or additional metadata may not be readily available. Further, it’s not easily scalable. We address these limitations via the novel task of summarizing news videos directly to entity-aware captions and establishing a large-scale dataset specifically for training and evaluation.

**External Knowledge.** To acquire information beyond the training set, many approaches leverage external knowledge sources such as Wikimedia (Chen et al., 2021), Google Search (Lin and Byrne, 2022), Wikidata (Hu et al., 2023), and so on. Orthogonally, Yu et al. (2023) argue that using an LLM to generate knowledge is better than retrieving it. Zhang et al. (2023b) extend this idea to images. Building on these works, we use LLMs as a knowledge source for videos. Unlike previous approaches, we propose a two-stage pipeline that combines learning and generation for more factual knowledge.

### 3 Dataset

In this section we describe our newly collected Video NEWS (VIEWS) dataset which contains




Figure 4: Word cloud of ‘PERSON’ entities in VIEWS.

news videos accompanied by captions and additional metadata such as the title, a paired news article and the news media source. We list further details below.

**Dataset Collection.** We build our dataset on top of an existing news video-article dataset, M<sup>2</sup>E<sup>2</sup>R (Ayyubi et al., 2022). M<sup>2</sup>E<sup>2</sup>R has a total of 220K videos-article pairs sourced from diverse news media outlets like Associated Press, BBC News, Axios, NPR etc. We leverage the articles paired with videos to generate caption for the videos. A straightforward approach would be to generate a summary from the article. However, articles often contain related historical stories, opinions, and predictions in addition to describing the video’s content. Including such information in the caption would render it unfaithful to the video, as this information is not visually depicted.

To mitigate this issue, we exploit an interesting property of M<sup>2</sup>E<sup>2</sup>R – some articles contain an explicit event description of each camera shot in the video, apart from the regular article story. We leverage these event descriptions and title (for context) to generate a video caption using a LLM (cf. Figure 2). Specifically, we use PaLM-2-L-IT (Anil et al., 2023) (more details in Appendix). This results in captions that are highly aligned with video content. The filtering step to collect videos having event descriptions in article leads to a total data size of ~144K. We split this data into 1.6K dev and test set each, leaving 141K for training.

**Quality Control and Annotation.** While the captions are already highly aligned with the video, we do further quality control, verification and rectification on dev/test set. We employ a human-in-the-loop procedure for this task. First, we ask a LLM rater to score the quality of captions according to the following criteria: 1) It should contain the name of important entities (people, place, organizations etc.); 2) It should not hallucinate ; 3) It should not leave out critical information from the event de-

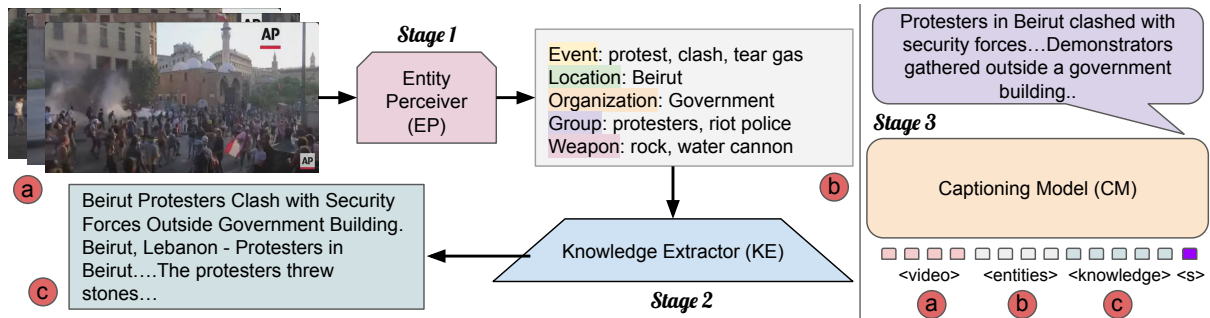


Figure 5: An overview of our proposed approach. Given the video input, we first use EP to detect entities. Then, KE uses the detected entities to extract contextual knowledge about the video. Finally, video, entities and context are input to CM to generate entity-aware video captions.

scriptions. The samples that the rater flagged as not being of high quality are separated for manual verification and correction. For rating, we used LLaMA-2 (Touvron et al., 2023).

Separately, we analysed how well the strategy of using LLM raters aligned with human judgement of a high quality caption. For this purpose, we manually rated a gold set of 100 samples and also asked the LLM raters to rate it. We found that only  $\sim 6\%$  times we flagged a caption to be of poor quality and LLM raters failed to flag it.

**Comparison with Existing Datasets.** As reported in Table 1, our VIEWS dataset significantly increases the scale of the news video captioning dataset by one to two orders of magnitude. Moreover, our VIEWS dataset is the first news video captioning dataset that has a high alignment between captions and videos as other datasets’ captions are derived from paired articles. This establishes the basis for training and evaluation of models generating captions from video content without access to paired articles. Compared to existing video datasets in other domains, our VIEWS dataset is focused on entity-aware captions, which cover numerous real-world entities, including people, places, events, organizations, etc. (see Figure 3 & Figure 4). The large-scale nature, the high video-caption alignment, and the focus on entity-aware caption make VIEWS first-of-its-kind in the community, to the best of our knowledge. We provide more dataset curation details, statistics, comparisons, and samples in Appendix A.

## 4 Method

Summarizing news videos directly to entity-aware captions is a challenging task. It requires not just recognition of named entities from people, places, and landmarks, but also an understanding of the context in which the events in the video are un-

folding. The straightforward way of directly training the model with supervision from entity-aware ground-truth captions may enable recognition of frequently occurring named entities. However, learning context this way is nearly impossible since context changes dynamically with every sample.

To address this, we divide the learning into three stages: 1) Entity Perceiver (EP) is trained to detect high-frequency entities from videos. 2) Knowledge Extractor (KE) uses detected entities to extract relevant context from external sources, eliminating the need for learning contexts. The extracted context also often contains missing long-tailed entities. 3) Captioning Model incorporates video, entities and context into informative entity-aware captions. The method is depicted visually in fig. 5. We detail each of the components below.

**Entity Perciever (EP).** The objective of EP is to recognize entities in the given video. Unlike in NLP, there doesn’t exist open-domain off-the-shelf entity detector from vision data. Consequently, we train our own Entity Perciever (EP). We use a transformer based sequence-to-sequence generative model instead of a discriminative model for EP. This choice is motivated by the fact that for open-domain entities the number of entity types (‘Person’, ‘Location’ etc.) and their instantiations may vary significantly across samples as shown in Figure 7. A generative model facilitates this property seamlessly. EP inputs videos and outputs entities relating each entity with its type in a dictionary format, as shown in fig. 5, Stage 1. To train EP, we extract ground-truth entities from the given caption label using a LLM. More details on this process in Appendix B.

Formally, we train EP  $\mathcal{F}_{EP} : \mathbb{R}^{H \times W \times 3 \times F} \rightarrow \mathbb{R}^{N \times L_E}$  using a language modelling objective (Sutskever et al., 2014) following (Wang et al.,



2022) with this loss,

$$L_{EP} = \frac{1}{L_E} \sum_i \text{CE}(\mathbf{e}_i, \mathcal{F}_{EP}(\mathbf{v}, \mathbf{e}_1, \dots, \mathbf{e}_{i-1})),$$

where  $L_E$  is the number of tokens in the target sequence, CE stands for the standard cross-entropy loss,  $\mathbf{v} \in \mathbb{R}^{H \times W \times 3 \times F}$  is the input video, and  $\mathbf{e}_i$  denotes the  $i$ -th ground-truth token in the target sequence. Note that this process is done separately before the captioning model is trained.

**Knowledge Extractor (KE).** KE is required to reason about the event in the video given entities and extract a relevant news data as context from an external knowledge source (depicted in Figure 5, Stage 2). In this work, we use a LLM for this purpose. LLMs are typically trained on massive corpus of internet text, which includes news data. This coupled with the LLM’s reasoning abilities allows extraction of relevant news for a given entities query. In contrast, KE sources such as Wikipedia (Hu et al., 2023) or Encyclopedia (Mensink et al., 2023) often don’t cover lesser known world affairs such as Bush’s tour of Africa Nation in Figure 1 or Cyril Svoboda and Condoleeza Rice’s meeting in Figure 7. Further, KE sources such as web searches lack LLM’s reasoning abilities and are quite sensitive to search query. Besides, there are limited public APIs available for web searches at this time.

Formally, this process entails generation of context sequence  $\mathbf{k}_1, \dots, \mathbf{k}_{L_K}$  by KE  $\mathcal{F}_{KE}$ , given a prompt sequence  $\mathbf{p}_1, \dots, \mathbf{p}_{L_P}$  and detected entities  $\hat{\mathbf{e}}_1, \dots, \hat{\mathbf{e}}_{L_E}$ . As shown in the equation:

$$\mathbf{k}_1, \dots, \mathbf{k}_{L_K} = \mathcal{F}_{KE}(\mathbf{p}_1, \dots, \mathbf{p}_{L_P}, \hat{\mathbf{e}}_1, \dots, \hat{\mathbf{e}}_{L_E}),$$

We detail the used prompt in Appendix B.

Notably, our approach keeps the LLM as an independent module and allows it to be replaced by any other sources of KE or even new/updated LLMs in future. We show how this enables our method to adapt to novel news content in Section 5.7.

**Captioning Model (CM).** Henceforth, we refer to the detected entities and extracted context together as Video Information ( $\mathcal{VI}$ ). CM integrates the video with  $\mathcal{VI}$  to generate entity and context aware video caption (illustrated in Figure 5, Stage 3). This is achieved via inbuilt transformer based attention mechanism inspired by (Lewis et al., 2020; Mensink et al., 2023). We demonstrate the effectiveness of our approach on three state-of-the-art CMs (more details in Section 5.1). It is important to note the versatility of our approach: it can be applied to any standard multimodal CM.

Model	B-4	R-L	CIDEr	Ent. F1
GIT	0.06	3.99	0.01	0.75
Video-LLaMA	0.40	12.90	0.80	1.45
GIT	5.40	22.06	18.04	16.20
+ $\mathcal{VI}$ (Ours)	5.80	22.45	22.27 $\Delta_{4.23}$	17.89 $\Delta_{1.69}$
BLIP-2	4.75	22.28	13.38	28.97
+ $\mathcal{VI}$ (Ours)	4.89	22.31	14.41 $\Delta_{1.03}$	29.11 $\Delta_{0.14}$
ViT-T5	3.05	18.29	7.24	19.20
+ $\mathcal{VI}$ (Ours)	3.48	18.92	9.87 $\Delta_{2.63}$	22.77 $\Delta_{3.57}$

Table 2: The impact of using our proposed approach (+ $\mathcal{VI}$ ) on generating entity-aware captions in VIEWS dataset. Our approach improves the performance of all models across all metrics. Greyed rows are evaluated in a zero-shot manner.

Despite the models’ different detailed architecture design, the training loss can be formalized as:

$$\mathcal{L}_{CM} = \frac{1}{L_C} \sum_i \text{CE}(\mathbf{h}_i, \mathcal{F}_{Cap}(\mathbf{v}, \hat{\mathbf{e}}_1, \dots, \hat{\mathbf{e}}_{L_E}, \mathbf{k}_1, \dots, \mathbf{k}_{L_K}, \mathbf{h}_1, \dots, \mathbf{h}_{i-1}))$$

where  $L_C$  is the length of caption,  $\mathcal{F}_{Cap}$  is the captioning model, and  $\mathbf{h}_i$  is  $i$ -th token in the ground-truth caption.

## 5 Experiments

### 5.1 Baselines

We carefully choose three baseline Captioning Models (CMs) to ensure we can analyse our approach on diverse paradigms of architectures/training strategies.

**GIT**(Wang et al., 2022) represents the class of traditional encoder-decoder model. We use the official pretrained GIT-L and finetune the model end-to-end with #frames=19, #maximum caption tokens=100, and #maximum  $\mathcal{VI}$  tokens=300.

**BLIP-2**(Li et al., 2023a) represents the category of models that exploit the power of (frozen) LLM via Q-Former (Li et al., 2023a) to generate captions from visual data. We finetune the official BLIP-2 pretrained weights with 13 frames, the maximum we can fit in the memory. The maximum caption and  $\mathcal{VI}$  tokens remain the same as GIT.

**ViT-T5.** GIT trains the model end-to-end but uses a relatively small decoder, while BLIP-2 uses strong LLM decoder but keeps it frozen. As such, we construct this baseline with a large LLM decoder (T5-L (Chung et al., 2022)) and finetune it end-to-end. The rest of the implementation details remain the same as BLIP-2.




Figure 6: Traditional captioning (MSRVTT) vs Entity-aware captioning (VIEWS). Performance drop  $>60\%$  on VIEWS across all models highlights its complexity.

Table 3: If available, video-paired article significantly boosts model performance. This implies unavailability of articles in VIEWS makes it challenging.

Method	CIDEr
GIT	17.78
+ Article	<b>22.20</b>

Additional baseline implementation details are provided in Appendix C. We report the main results on the test set. All ablations and analyses are reported on the validation set using GIT, unless otherwise stated.

## 5.2 Experimental Setup

**Implementation Details.** For the Entity Perceiver (EP), we use a pre-trained GIT model and for Knowledge Extractor (KE), we use PaLM-2-L-IT LLM (Anil et al., 2023).

**Training Details.** We train GIT, BLIP-2 and ViT-T5 for 42, 40 and 30 epochs respectively, using a batch size of 128, a learning rate of  $1e-5$ , adamW (Loshchilov and Hutter, 2017) optimizer and cosine scheduler with a linear warmup.

**Metrics.** We use the standard BLEU-4 (B-4), Rouge-L (R-L) and CIDEr metrics used in captioning tasks (Vedantam et al., 2015; Lin, 2004; Papineni et al., 2002). We also calculate an Entity  $F_1$  (Ent.  $F_1$ ) score to measure how well the ground-truth entities are covered in the generated caption, similar to (Liu et al., 2021; Bertsch et al., 2023). To this end, named-entities are extracted from ground-truth caption and generated caption, and then compared using  $F_1$  score. Among all metrics, we emphasize CIDEr and Ent  $F_1$  the most. CIDEr is important for evaluating caption quality as it strikes a balance between precision focused B-4 and recall focused R-L. On the other hand, Ent  $F_1$  is crucial for evaluating entities presence in the captions.

Ablation	CIDEr
1. Ours	<b>21.00</b>
2. - w/o Entities	20.31
3. - w/o Knowledge	20.75
4. - w/o $\mathcal{V}\mathcal{I}$	17.78

Table 4: Ablation study on GIT to study the impact of our design choices. Both entities and context positively impact the model’s performance.

GT Caption Source	CIDEr
Event Descriptions	<b>18.04</b>
Paired Article	10.5

Table 5: Ground-Truth (GT) captions sourced from articles perform poorly compared to captions from event descriptions. This indicates captions derived from articles are not well aligned with the video

## 5.3 Main Results: Proposed Method’s Impact

We study how incorporating entities and context ( $\mathcal{V}\mathcal{I}$ ) to the videos via our proposed approach (+ $\mathcal{V}\mathcal{I}$ ) improves CMs’ performance on the task of generating entity-aware captions. The results are reported in Table 2. We make the following observation:

**Our method improves all three captioning models’ performance.** For all three models, our approach registers improvements across all the metrics with significant increments in CIDEr (Avg.  $\sim 22.5\%$ ) and Ent.  $F_1$  (Avg.  $\sim 10\%$ ). This demonstrates our approach’s effectiveness in infusing existing CMs with the ability to generate context and entity aware captions. Moreover, it signifies our method’s model-agnostic behavior.

**End-to-end finetuned models improve most.** Avg. improvement for finetuned models (GIT and ViT-T5) is much larger when compared to partially-tuned BLIP-2 (30% vs 8% CIDEr and 14.5% vs 0.5% Ent.  $F_1$ ). This indicates end-to-end training is best suited to adapt models to utilize our VI.

**Zero-shot models perform poorly.** Git and even LLM based, Video-LLaMA, perform quite poorly on this task. This highlights the complexity of the task and exposes existing model’s inability to generate entity-aware captions.

## 5.4 Analysis of Task Complexity

The low zero-shot and even finetuned baseline performance on VIEWS hints at high task complexity. In this section, we analyze various factors that make it so.

**Entity-aware video captioning is challenging.** In Figure 6, we compare CIDEr scores of three

baselines on a traditional captioning benchmark, MSRVT, against our entity-aware benchmark, VIEWS. We find that all the baselines undergo at least 61% performance drop from MSRVT to VIEWS. This emphasizes the complexity of our task.

**Unavailability of video paired article makes VIEWS challenging.** VIEWS requires generating entity-aware captions directly from videos to be applicable in practical settings where video-paired news article is unavailable. If available, we find that it leads to performance boost of 24.85% in table Table 3. This illustrates that part of VIEWS’s complexity comes from unavailability of paired article.

**Captions derived from paired article are not well aligned with the video and thus noisy.** To justify our usage of event descriptions instead of paired articles to generate ground-truth video captions, we train the model on captions generated from event descriptions and paired articles respectively. We report the results in Table 5. We find that training the model on article-sourced-captions leads to a drop of 71.81% performance as compared to event-descriptions-sourced-captions. This demonstrates that the paired article, and hence article-sourced-caption, has a noisy relationship with the visual content.

### 5.5 Analysis of Proposed Method

The proposed method relies heavily on detected entities and extracted context. As such, we evaluate their quality and impact on model performance.

**Entities are critical.** We evaluate the quality of our detected entities in Table 6. For reference, we use a zero-shot CLIP (Radford et al., 2021) model. We find that our entities are 50× better. Qualitative analysis in Figure 7 shows that our approach is quite good in recognizing people (Condoleezza Rice in top row), location (Johannesburg and Iran in middle and bottom row respectively), and diverse events (political meeting, cultural event, elections etc). It does this from visual data alone, implying that it can learn to recognize places from landmarks and cultural attire, people from a few samples, and organizations from people and locations.

When entities are removed from our pipeline, the model performance drops (see Table 4). Figure 8 shows that the model performance is directly proportional to the quality of entities. All this emphasizes the importance of entities. We further note that there’s a large performance gap between

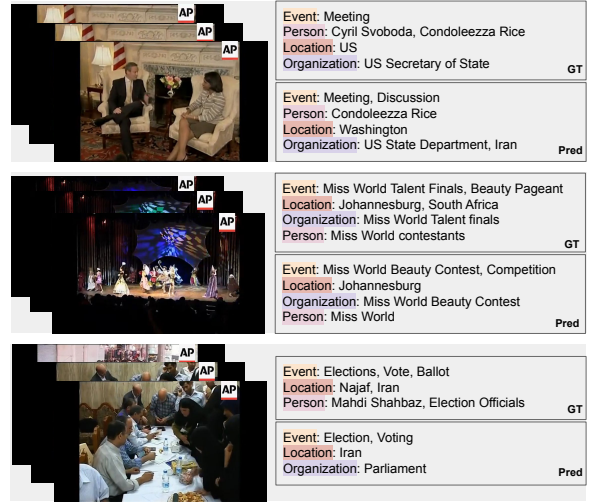


Figure 7: Qualitative analysis of detected entities. Our approach can recognize even obscure entities like Johannesburg, Iran, and so on.




Table 6: Quantitative analysis of detected entities vs CLIP.

Figure 8: Model performance (CIDEr) improves linearly with entities (F<sub>1</sub>).

detected entity (Ent. F<sub>1</sub>=10) and ground-truth entity (Ent. F<sub>1</sub>=1), indicating large returns on investment in entity-perception development. We provide more details and analysis on entities in appendix G.

**Context matters.** To assess the quality of our context, we compare it to the context generated by different versions of PaLM-2 and LLaMA-2 (Touvron et al., 2023) as shown in Table 7. We measure context quality using BERTScore (BSc), measuring contextual relevance with the ground-truth caption, as traditional metrics (B-4, R-L, and CIDEr) are unsuitable for long paragraph comparisons. We observe that our context scores higher than the others. As expected, a general trend is that larger LLMs achieve higher scores.

The importance of context in our pipeline can be deduced from Table 4: the performance drops if it’s removed. Further, Figure 10 shows that model performance improves exponentially with higher context BSc. This implies that even minor improvements in context quality can drive significant gains in model performance, highlighting a promising area for future research.

**Single stage vs our two-stage KE.** Video-based LLMs, like Video-LLaMA (Zhang et al., 2023a) and VideoChat (Li et al., 2024), have proven to




Figure 9: Comparison between our method and a video-based LLM (Video-LLaMA) to extract context. Video-LLaMA fails to extract relevant news knowledge; rather it just describes the video content in a generic manner.

KE LLM	BSc
PaLM-2-340B	<b>85.54</b>
PaLM-2-24B	85.50
LlAMA-2-13B	85.09
LlAMA-2-7B	83.56

Table 7: LLMs ablation as KE. Ours is PaLM-2-340B.




Figure 10: Model performance (CIDEr) improves exponentially with context (BSc).

Method	Bert-Score	Ent. F <sub>1</sub>
Single-Stage KE (Video-LLaMA)	71.11	1.24
Two-Stage KE (Ours)	<b>85.54</b>	<b>13.53</b>

Table 8: Our two-stage KE significantly outperforms Video-LLaMA’s single-stage KE, showing that extracting context from detected entities is better than from videos directly.

be effective in understanding visual content and leveraging related factual knowledge. As such, we strive to study if a Video-LLaMA based single-stage KE can simply replace our two-stage pipeline of first learning entities to then extract contextual knowledge. To this end, we input the video to Video-LLaMA and task it to extract a relevant news article. We observe in Table 8 that Video-LLaMA glaringly underperforms our approach despite exhaustive prompt engineering.

Qualitative comparison of caption, Video-LLaMA context, and our context in Figure 9 shows that Video-LLaMA keeps trying to describe the video instead of extracting a relevant news information. In comparison, the context extracted by our method is quite good: top row shows that we can extract information about the trial of Michael Jackson’s doctor, Conrad Murray, just from the visual data. This demonstrates that extracting context from learned entities, rather than directly from videos, improves factuality by reducing task complexity. It also implies that, despite the success of video-based LLMs, they still lag behind text-based LLMs on knowledge-intensive tasks like ours.



Figure 11: Qualitative results of our Git +  $\mathcal{V}$ Z model on VIEWS. Green and red highlights denote correctly and wrongly detected entities respectively.

## 5.6 Qualitative Analysis

We qualitatively analyse the predictions from our approach in Figure 11. We see that our method effectively summarizes news video while being informative by predicting relevant named-entities (Nawaz Sharif, Gazprom, Ukraine etc). Remarkably, it uses the contextual information to even detect entities not mentioned in the ground-truth captions: Lahore in row 1. While it sometimes predicts a wrong entity (Russia Deputy Prime Minister in row 2), it is able to generally summarize video content well (gas crisis in Ukraine in row 2). These qualitative results further provide evidences of our proposed approach’s effectiveness.

We provide more qualitative samples and analysis of failure cases in Appendix F.

## 5.7 Generalization.

**Generalization over time.** Our use of an LLM as a plug-and-play Knowledge Extractor (KE) offers unique advantages. It allows us to replace our current LLM with a better or more updated one, possessing more recent knowledge or even web search capabilities, while keeping the rest of the model unchanged. This enables our approach to generalize over time to novel news content. The only caveat is that our frozen CM should adapt to new contexts without re-training. To test this, we



Model	B-4	R-L	CIDEr	Ent. F1
GIT	4.59	19.58	16.51	19.86
+ $\mathcal{V}\mathcal{I}$ (Ours)	4.65	19.60	18.01 $\Delta$ 1.50	20.72 $\Delta$ 0.96

Table 9: Our method generalizes across time when trained on data before 1 Jan, 2017 and evaluated on data after that date. This demonstrates our method’s ability to adapt to new context when available.

trained the model on data up to January 1, 2017, and the data from after that date acts as new content for evaluation. Since our LLM’s cut-off date is beyond content date in eval set, it suffices as “updated” LLM. The reported results in Table 9 demonstrate that our approach can indeed adapt to novel content, and thus generalize over time.

**Generalization across datasets and modality.** As discussed before, existing video news datasets contain captions containing information from articles as well, which makes them unsuitable for direct video to caption generation. However, we test the limits of our proposed approach on one such dataset, MM-AVS (Fu et al., 2021) in Appendix H. We also evaluate our approach on an news image captioning dataset (Liu et al., 2021) in Appendix H. We find that our approach is able to generalize to these new benchmarks as well.

## 6 Conclusion

In this work, we proposed the new task of generating entity-aware captions from videos only. We also introduced a dataset, VIEWS, to help spur research in this field. The task is challenging for existing captioning models as it requires recognizing named-entities in the video and also the context in which the event in video is happening. We proposed a modular approach to address these challenges, which we demonstrate improves over vanilla video-only model. Still, the relatively modest numbers suggest the task remains challenging and exciting for further research. We provide insights into these future research directions through our rich experiments and ablation studies

## Limitations

We present a benchmark that challenges models to detect entities and context from visual cues, and use them to generate entity-rich video captions. We assume that there’s a knowledge source out there that does contain information about this event. We focus the evaluation of this benchmark to detect correct entities, extract relevant context, and in-

tegrate them into a coherent caption. We do not cater to videos beyond the scope of this dataset or daily news videos. While this is a nice property to have, our evaluation does not prioritize this, but our method design does allow modular updates with new KEs to take care of these situations. We leave this for future work.

## Acknowledgements

We want to thank Junzhang Liu and Venkat Suprabath Bitra for helping us with the human hallucination study.

## References

- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, et al. 2022. [Flamingo: a visual language model for few-shot learning](#). *ArXiv preprint*, abs/2204.14198.
- Rohan Anil, Andrew M. Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, et al. 2023. Palm 2 technical report. In *arXiv preprint arXiv:2305.10403*.
- Hammad A. Ayyubi, Christopher Thomas, Lovish Chum, Rahul Lokesh, Yulei Niu, Xudong Lin, Long Chen, Jaywon Koo, Sounak Ray, and Shih-Fu Chang. 2022. [Multimodal event graphs: Towards event centric understanding of multimodal world](#). *Preprint*, arXiv:2206.07207.
- M. Bain, A. Nagrani, A. Brown, and A. Zisserman. 2021. [Condensed movies: story based retrieval with contextual embeddings](#). *Computer Vision – ACCV 2020*, pages 460–479.
- Amanda Bertsch, Uri Alon, Graham Neubig, and Matthew R. Gormley. 2023. [Unlimiformer: Long-range transformers with unlimited length input](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Ali Furkan Biten, Lluís Gomez, Marçal Rusiñol, and Dimosthenis Karatzas. 2019. [Good news, everyone! context driven entity-aware captioning for news images](#). *2019 IEEE CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12458–12467.
- David Chen and William B Dolan. 2011. Collecting highly parallel data for paraphrase evaluation. In *Proceedings of the 49th annual meeting of the association for computational linguistics: human language technologies*, pages 190–200.
- Kezhen Chen, Qiuyuan Huang, Daniel McDuff, Yonatan Bisk, and Jianfeng Gao. 2021. Krit: Knowledge-reasoning intelligence in vision-language transformer. Technical Report MSR-TR-2021-34, Microsoft.

- Xi Chen, Josip Djolonga, Piotr Padlewski, Basil Mustafa, Soravit Changpinyo, Jialin Wu, Carlos Riquelme Ruiz, Sebastian Goodman, Xiao Wang, Yi Tay, Siamak Shakeri, Mostafa Dehghani, Daniel M. Salz, Mario Lucic, Michael Tschannen, Arsha Nagrani, Hexiang Hu, Mandar Joshi, Bo Pang, Ceslee Montgomery, Paulina Pietrzyk, Marvin Ritter, A. J. Piergiovanni, Matthias Minderer, Filip Pavetic, Austin Waters, Gang Li, Ibrahim M. Alabdulmohsin, Lucas Beyer, Julien Amelot, Kenton Lee, Andreas Steiner, Yang Li, Daniel Keysers, Anurag Arnab, Yuanzhong Xu, Keran Rong, Alexander Kolesnikov, Mojtaba Seyedhosseini, Anelia Angelova, Xiaohua Zhai, Neil Houlsby, and Radu Soricut. 2023. **Pali-x: On scaling up a multilingual vision and language model**. *ArXiv*, abs/2305.18565.
- Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, et al. 2022. **Pali: A jointly-scaled multilingual language-image model**. *arXiv preprint arXiv:2209.06794*.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. **Scaling instruction-finetuned language models**. *arXiv preprint arXiv:2210.11416*.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. **An image is worth 16x16 words: Transformers for image recognition at scale**. *arXiv preprint arXiv:2010.11929*.
- Yuxin Fang, Quan Sun, Xinggang Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. 2023. **Eva-02: A visual representation for neon genesis**. *arXiv preprint arXiv:2303.11331*.
- Xingyu Fu, Ben Zhou, Ishaan Chandratreya, Carl Vondrick, and Dan Roth. 2022. **There’s a time and place for reasoning beyond the image**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1138–1149, Dublin, Ireland. Association for Computational Linguistics.
- Xiyan Fu, Jun Wang, and Zhenglu Yang. 2021. **MM-AVS: A full-scale dataset for multi-modal summarization**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5922–5926, Online. Association for Computational Linguistics.
- Ziniu Hu, Ahmet Iscen, Chen Sun, Zirui Wang, Kai-Wei Chang, Yizhou Sun, Cordelia Schmid, David A. Ross, and Alireza Fathi. 2023. **Reveal: Retrieval-augmented visual-language pre-training with multi-source multimodal knowledge memory**. *2023 IEEE CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 23369–23379.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. **Survey of hallucination in natural language generation**. *ACM Computing Surveys*, 55(12):1–38.
- Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. 2017. **Dense-captioning events in videos**. *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 706–715.
- Mateusz Krubiński and Pavel Pecina. 2023. **MLASK: Multimodal summarization of video-based news articles**. In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 910–924, Dubrovnik, Croatia. Association for Computational Linguistics.
- Jie Lei, Licheng Yu, Tamara L. Berg, and Mohit Bansal. 2020. **Tvr: A large-scale dataset for video-subtitle moment retrieval**. In *European Conference on Computer Vision*.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. **Retrieval-augmented generation for knowledge-intensive nlp tasks**. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven C. H. Hoi. 2023a. **Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models**. In *International Conference on Machine Learning*.
- KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. 2024. **Videochat: Chat-centric video understanding**. *Preprint*, arXiv:2305.06355.
- Mingzhe Li, Xiuying Chen, Shen Gao, Zhangming Chan, Dongyan Zhao, and Rui Yan. 2020. **VMSSMO: Learning to generate multimodal summary for video-based news articles**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9360–9369, Online. Association for Computational Linguistics.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. 2023b. **Evaluating object hallucination in large vision-language models**. In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- Yunlong Liang, Fandong Meng, Jinan Xu, Jiaan Wang, Yufeng Chen, and Jie Zhou. 2023. **Summary-oriented vision modeling for multimodal abstractive summarization**. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2934–2951, Toronto, Canada. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. **Rouge: A package for automatic evaluation of summaries**. In *Text summarization branches out*, pages 74–81.

- Weizhe Lin and Bill Byrne. 2022. [Retrieval augmented visual question answering with outside knowledge](#). *Preprint*, arXiv:2210.03809.
- Fuxiao Liu, Yinghan Wang, Tianlu Wang, and Vicente Ordonez. 2021. [Visual news: Benchmark and challenges in news image captioning](#). *Preprint*, arXiv:2010.03743.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Di Lu, Spencer Whitehead, Lifu Huang, Heng Ji, and Shih-Fu Chang. 2018. [Entity-aware image caption generation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4013–4023, Brussels, Belgium. Association for Computational Linguistics.
- Thomas Mensink, Jasper Uijlings, Lluís Castrejon, Arushi Goel, Felipe Cadar, Howard Zhou, Fei Sha, André Araujo, and Vittorio Ferrari. 2023. Encyclopedic vqa: Visual questions about detailed properties of fine-grained categories. *arXiv preprint arXiv:2306.09224*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Jielin Qiu, Jiacheng Zhu, Mengdi Xu, Franck Dernoncourt, Trung Bui, Zhaowen Wang, Bo Li, Ding Zhao, and Hailin Jin. 2022. [Semantics-consistent cross-domain summarization via optimal transport alignment](#). *Preprint*, arXiv:2210.04722.
- 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. 2021. [Learning transferable visual models from natural language supervision](#). *Preprint*, arXiv:2103.00020.
- Anna Rohrbach, Atousa Torabi, Marcus Rohrbach, Niket Tandon, Christopher Pal, Hugo Larochelle, Aaron Courville, and Bernt Schiele. 2016. [Movie description](#). *Preprint*, arXiv:1605.03705.
- Marcus Rohrbach, Anna Rohrbach, Michaela Regneri, Sikandar Amin, Mykhaylo Andriluka, Manfred Pinkal, and Bernt Schiele. 2015. [Recognizing fine-grained and composite activities using hand-centric features and script data](#). *International Journal of Computer Vision*, pages 1–28.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27.
- Peggy Tang, Kun Hu, Lei Zhang, Jiebo Luo, and Zhiyong Wang. 2022. [Tldw: Extreme multi-modal summarisation of news videos](#). *Preprint*, arXiv:2210.08481.
- OpenAI Team. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *Preprint*, arXiv:2307.09288.
- Alasdair Tran, Alexander Mathews, and Lexing Xie. 2020. Transform and tell: Entity-aware news image captioning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575.
- Jianfeng Wang, Zhengyuan Yang, Xiaowei Hu, Linjie Li, Kevin Lin, Zhe Gan, Zicheng Liu, Ce Liu, and Lijuan Wang. 2022. [Git: A generative image-to-text transformer for vision and language](#). *Preprint*, arXiv:2205.14100.
- Shuhe Wang, Xiaofei Sun, Xiaoya Li, Rongbin Ouyang, Fei Wu, Tianwei Zhang, Jiwei Li, and Guoyin Wang. 2023. [Gpt-ner: Named entity recognition via large language models](#). *Preprint*, arXiv:2304.10428.
- Spencer Whitehead, Heng Ji, Mohit Bansal, Shih-Fu Chang, and Clare Voss. 2018. [Incorporating background knowledge into video description generation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3992–4001, Brussels, Belgium. Association for Computational Linguistics.
- Jun Xu, Tao Mei, Ting Yao, and Yong Rui. 2016. [Msr-vid: A large video description dataset for bridging](#)

- video and language. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5288–5296.
- Antoine Yang, Arsha Nagrani, Paul Hongsuck Seo, Antoine Miech, Jordi Pont-Tuset, Ivan Laptev, Josef Sivic, and Cordelia Schmid. 2023. [Vid2seq: Large-scale pretraining of a visual language model for dense video captioning](#). *Preprint*, arXiv:2302.14115.
- Wenhao Yu, Dan Iter, Shuohang Wang, Yichong Xu, Mingxuan Ju, Soumya Sanyal, Chenguang Zhu, Michael Zeng, and Meng Jiang. 2023. Generate rather than retrieve: Large language models are strong context generators. In *International Conference for Learning Representation (ICLR)*.
- Hang Zhang, Xin Li, and Lidong Bing. 2023a. [Video-llama: An instruction-tuned audio-visual language model for video understanding](#). *arXiv preprint arXiv:2306.02858*.
- Zhuosheng Zhang, Aston Zhang, Mu Li, Hai Zhao, George Karypis, and Alex Smola. 2023b. Multi-modal chain-of-thought reasoning in language models. *arXiv preprint arXiv:2302.00923*.
- Sanqiang Zhao, Piyush Sharma, Tomer Levinboim, and Radu Soricut. 2019. [Informative image captioning with external sources of information](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6485–6494, Florence, Italy. Association for Computational Linguistics.
- Luowei Zhou, Chenliang Xu, and Jason J Corso. 2018. [Towards automatic learning of procedures from web instructional videos](#). In *AAAI Conference on Artificial Intelligence*.
- Mingyang Zhou, Grace Luo, Anna Rohrbach, and Zhou Yu. 2022. [Focus! relevant and sufficient context selection for news image captioning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 6078–6088, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.



## Appendix

We provide additional details and analysis to further support the contributions of this work:

- Dataset Details (Appendix A)
- Method Details (Appendix B)
- Baselines Details (Appendix C)
- Effect of Predicted Entity Structure (Appendix D)
- Effect of utilizing ASR (Appendix E)
- Qualitative Examples (Appendix F)
- Further Analysis of Detected Entities (Appendix G)
- Generalization across Dataset and Modality (Appendix H)
- Case Study on Human Performance on the Dataset (Appendix I)

### A Dataset Details

We provide more samples from our dataset in Figure 12. The samples highlight key properties of our dataset:

- The bullet summaries used to generate ground-truth captions are highly video-aligned and precise event descriptions of the video.
- The ground-truth captions are very faithful to the video content.
- The ground-truth captions are rich in entities. Figure 13 plots the distribution of entities per caption demonstrating that majority of the captions contain around 4-5 entities.

In Figure 14, we illustrate our dataset’s high video-caption alignment compared to prior News Video Captioning datasets (MM-AVS and VMSMO). The source of information for each sentence in the caption is depicted by the color of the sentence. From the figure, we see that captions in both MM-AVS and VMSMO are sourced from videos *and* paired articles. This implies that video is not enough to generate ground-truth captions. In comparison, all sentences in VIEWS’s captions are sourced from video. This leads to high video-caption alignment, allowing captions to be generated directly from videos without needing a

paired article. Note that MM-AVS and VMSMO represent the general issue of low video-caption alignment in all existing News Video Captioning datasets.

The reason for high video-caption alignment in our dataset is rooted in our pipeline for generating ground-truth captions. We utilize the creator-provided video scene descriptions, uniquely available in our parent dataset (M<sup>2</sup>E<sup>2</sup>R), to generate video captions. This leads to captions whose information source is grounded in videos. This property gives our dataset a distinctive advantage over other news video captioning datasets. The full comparison to prior datasets is presented in Table 10.

#### LLM Based Ground-Truth Caption Generation.

The first step in generating ground-truth video captions is to filter out the highly video-aligned and precise event descriptions (bullet summaries) from the paired article. We use simple regex for this purpose. The next step is to generate captions from the bullet summaries. Towards this goal, we utilize PaLM-2-L-IT (Anil et al., 2023) large language model (LLM). We input to the LLM the following prompt:

```
Task: In this task, you will be given a context and your task is to summarize it in a few sentences. Please don't include specific dates in the summary.
```

```
Context: <bullet_summaries >
```

```
Instruction: Generate a summary of the Context in around 40 words.
```

The resulting captions are highly aligned with the video content as they have been summarized from video event descriptions provided in the paired article. We also specifically remove dates from captions as dates are quite difficult to predict from just the visual signal (Fu et al., 2022).

**LLM Based Caption Quality Control.** While the generated captions are already of high quality, we further control the quality of captions for test and dev set. To this end, we input the captions to LLaMA-2 LLM ask it to rate whether the captions is of high quality or not. We provide the LLM the following prompt for this purpose:

```
In this task, you will be given a context and its summary. Your job is to determine whether the quality of the summary is high or not.
```

```
A high quality summary has all of the following characteristics:
```

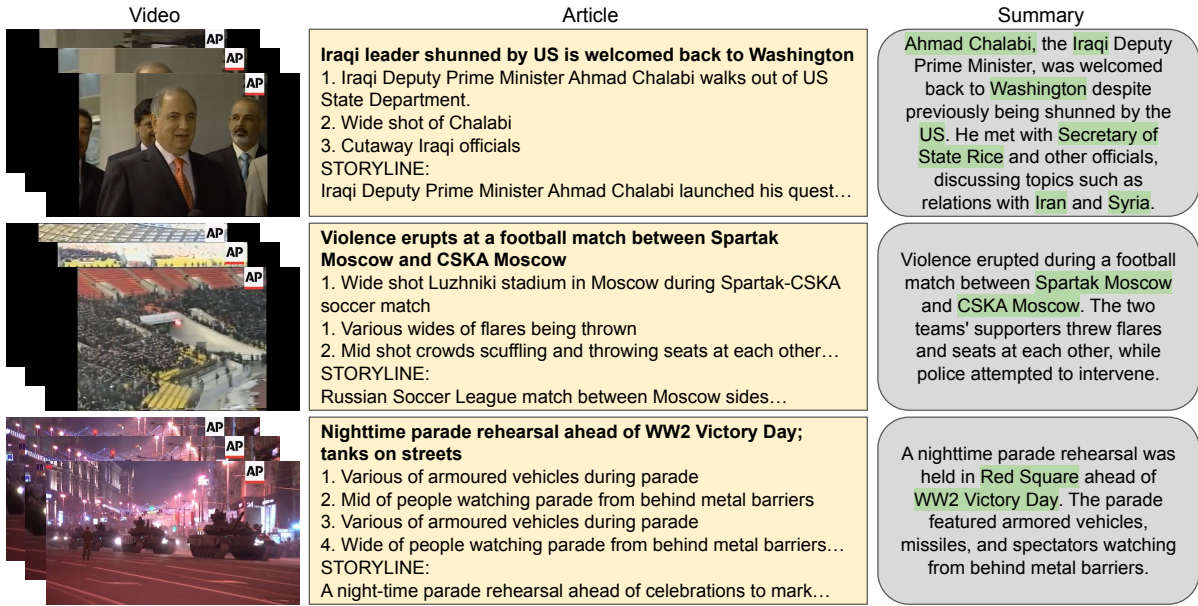


Figure 12: We present more samples from our dataset highlighting key properties: high caption alignment with videos and rich presence of entities (marked in green).

Dataset	Domain	#Videos	#Length Avg. (s)	#Words Avg.	Entities in Cap.	Cap.-Vid Align ↑
MSR-VTT (Xu et al., 2016)	Open	7.18K	20	13	✗	✓
ActyNet Cap (Krishna et al., 2017)	Open	20K	180	14	✗	✓
YouCook II (Zhou et al., 2018)	Cooking	2K	316	-	✗	✓
TVR/TVC (Lei et al., 2020)	TV show	21.8K	76.2	13.4	✗*	✓
LSMDC (Rohrbach et al., 2016)	Movie	128K	4.1	9	✗*	✓
MM-AVS (Fu et al., 2021)	News	2.17K	109	57	✓	✗
XMSMO (Tang et al., 2022)	News	4.8K	346	12	✓	✗
MLASK (Krubiński and Pecina, 2023)	News	41.2K	86	33	✓	✗
VIEWS (Ours)	News	144K	148.9	46	✓	✓

Table 10: Additional details on dataset comparison.




Figure 13: Entity distribution per caption.

- It contains the important entities ( person, place, organization etc.) in the context
- It does not contain hallucinations
- It does not leave out critical information from the context

Context: <bullet\_summaries >

Summary: <summary>

Instruction: Is Summary of high quality?  
Please choose between Yes and No.

As discussed in section 3, the captions which are flagged poor is sent for manual verification and correction. This has been done by the authors themselves. In addition to LLaMA-2, we also used PaLM-2 to rate the captions and found similar results.

## B Method Details

**Ground-Truth Entity Extraction.** We detect entities from video using Entity Perceiver (EP). To train EP, we require ground-truth entities. This is generated by extracting entities from bullet summaries using a LLM (PaLM-2-L-IT). We use a LLM for extracting these ground-truth entities as they have shown to excel in these kinds of Infor-

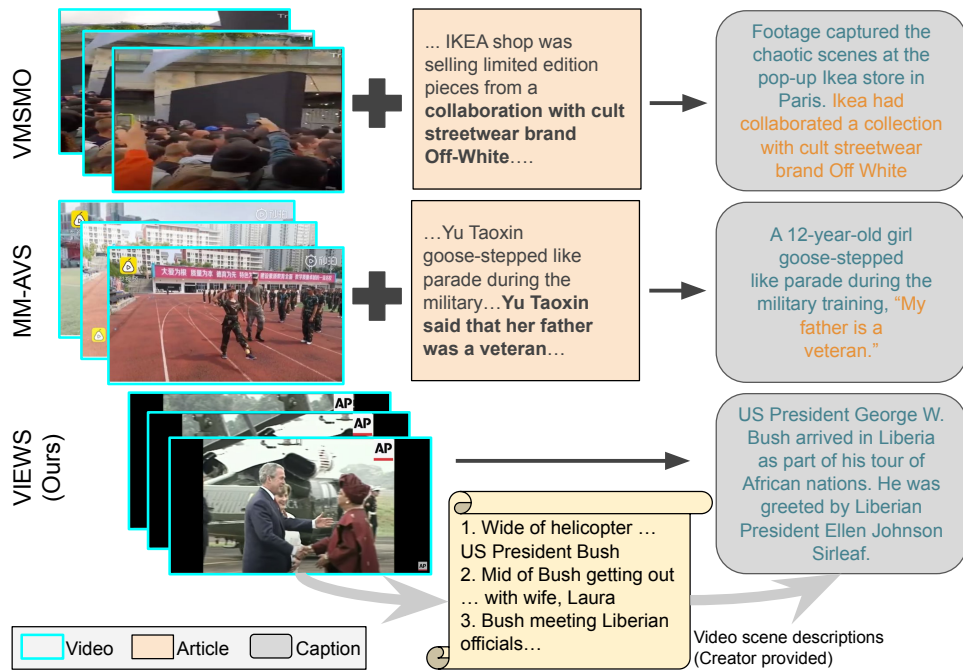


Figure 14: Video-Caption alignment comparison of VIEWS (Ours) against prior News Video Captioning datasets. For each dataset, sentences in caption are colored according to their source: video or article. MM-AVS (Top) and VMSMO (Middle) both contain sentences in captions that are sourced from articles; thus their captions are not well-aligned with the video. In contrast, all sentences in VIEWS are sourced from video, resulting in high video-caption alignment. The key to this high alignment is generating captions from creator provided video scene descriptions.

mation Extraction (IE) Named-Entity Recognition (NER) tasks (Wang et al., 2023). To this end, we input the following prompt to the LLM:

Task: In this task, you will be provided a piece of text and your task would be to extract named entities from the text. Please don't include any date entity.

Text: <bullet\_summaries >

Instruction: Extract named entities from the following text in this format: {<Entity\_Type >: [<Entity\_list >]}.

Again, we don't extract any date entity as they are difficult to predict from visual data.

**Knowledge Extraction (KE).** Once we predict entities using KP, we utilize them to extract a relevant news article as a knowledge source. We again use the same LLM as KE for this purpose. The prompt given is:

Task: In this task, you will be provided a list of entities and your task would be to retrieve a summary of the most relevant news article related to those entities.

Entities: <entities >

Instruction: Retrieve a summary of the

most relevant news article involving these Entities.

While use of online web search may seem a natural choice for KE, they have certain limitations. They are sensitive to search query and can be quite noisy. Besides, only limited search APIs are available publicly. For example, Google search charges for > 100 queries/day<sup>2</sup>.

## C Baselines Details

**GIT**(Wang et al., 2022) is a state-of-the-art vision-language model that has demonstrated strong results on 12 benchmarks by scaling up data and model size. To finetune it on our task, we follow the proposed approach of concatenating all frames' patch embeddings together and feeding them to the model to generate captions. We also concatenate frames' tokens with  $\mathcal{V}$  tokens before inputting to the model. We use the official pretrained GIT-L with ViT-L/14 image encoder(Dosovitskiy et al., 2020) and 6 layer transformer decoder (Vaswani et al., 2017).

**BLIP-2**(Li et al., 2023a) has shown strong results on many knowledge-intensive tasks by leveraging

<sup>2</sup><https://developers.google.com/custom-search/v1/overview?hl=ja>

a frozen LLM in its architecture. They utilize a Q-Former (Li et al., 2023a) to bridge the embedding space between visual signals and textual space of a LLM. During training, only Q-Former is learned while keeping LLM and the image encoder frozen. We also finetune only the Q-Former as in the original implementation. We utilize the Q-Former to extract queries for each frames, concatenate them with temporal embeddings, and then join them together before feeding into the LLM. Again, for our pipeline, these query embeddings are also concatenated with context tokens before inputting to LLM. We finetune the official BLIP-2 pretrained weights obtained using ViT-L/14 image encoder and Flan-T5-XL LLM (Chung et al., 2022).

**ViT-T5.** We construct this baseline to leverage the best of both LLMs and fine-tuning. To this end, ViT-g/14 (Fang et al., 2023) is employed as the image encoder and T5-L (Chung et al., 2022) as the LLM decoder. Per frame Q-Former is utilized, as in BLIP-2, to extract query embeddings and appended with temporal embeddings. These query embeddings together with our entity and context tokens are concatenated together to form the input to the LLM decoder. We finetune it end-to-end. The model is pretrained on WebLI data (Chen et al., 2022) and scores a strong CIDEr (Vedantam et al., 2015) of 69.2 on MSRVT Video Captioning benchmark when finetuned.

## D Effect of Predicted Entity Structure

We predict entities in a structured format as shown in fig. 7. This structure is in the form of a dictionary and follows this format:

$$\{ \langle Entity\_Type \rangle : [ \langle Entity\_List \rangle ] \}$$

The structure relates each entity to its type. This information is important in the downstream task of Knowledge Extraction (KE). For example, in the case of a United Nations (UN) summit held at, say, Geneva, the United States (US) acts as an organization not a location. Not delineating the entity type of United States may give the false impression that the UN summit was held in the US. We measure the quantitative impact of this structured entity detection in table 11. We find that there is marked degradation in the quality of extracted knowledge (Ent. F<sub>1</sub> 13.53 to 12.32), if we don't input entities in the structured format. This demonstrates the usefulness of detecting entities in the aforementioned format.

Model	Bert-Score	Ent. F <sub>1</sub>
Ours	<b>85.54</b>	<b>13.53</b>
- w/o Entities Structure	85.47	12.32

Table 11: Comparison of extracted knowledge with and without entity structure.

Model	BLEU-4	Rouge-L	CIDEr
GIT	5.41	21.88	17.78
+ ASR	5.42	21.91	17.84
GIT + Context	<b>5.62</b>	<b>22.12</b>	<b>21.0</b>
+ ASR	5.57	22.0	20.22

Table 12: Effect of using ASR on our task.

## E Effect of utilizing ASR

Automatic Speech Recognition (ASR) is an important tool in video captioning. It captures audio signals from the video in text form and provides this information as complementary to visual signals. This information is especially impactful in summarizing instructional videos (Yang et al., 2023). It would also have been beneficial in a narrative or reportive news video. However, most videos in our dataset is descriptive, in the sense that it just captures the news event video without any narrator or reporter. This is exemplified by the fact that only 40% of our data contains any ASR content. To further illustrate the insignificant impact of ASR on our dataset, we compare our models with a version trained with ASR signals. We report the results in table 12. We see that adding ASR to GIT barely improves the performance (17.84 vs 17.78 CIDEr). This is far below the improvement recorded by our approach which adds entities and external knowledge as Context (21.0 vs 17.78 CIDEr). We further note that adding ASR to our Context depreciates model's performance a bit (21.0 to 20.22 CIDEr). We hypothesize that this may be due to the fact that our approach already provides relevant context to the model and adding ASR may just create noise. These results highlight that while ASR may be useful for some video captioning tasks, it's not very relevant to our news summarization task.

## F Qualitative Examples

We provide additional qualitative examples in fig. 15. These examples further illustrate our approach's effectiveness in generating entity-aware captions. The model is able to detect people



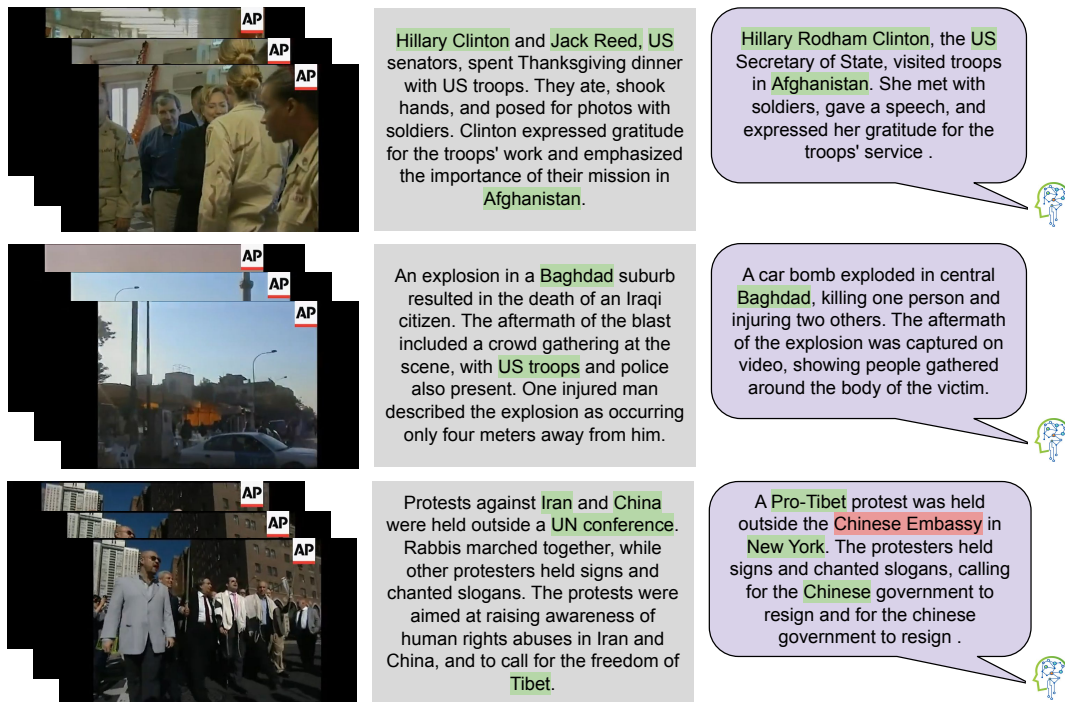


Figure 15: Additional sample predictions from our approach VIEWS. Left: Video; Middle: Ground-Truth Captions; Right: Prediction. Green highlights denote correctly detected entities, while red highlights denote wrongly predicted entities.

(Hillary Clinton in row 1), place (Afghanistan in row 1 and Baghdad in row 2) and even detect novel entity not present in ground-truth (New York in row 3). Further, the model is able to effectively integrate these entities into informative captions.

**Failure Analysis.** While our method generally performs well, we seek to analyze its failure cases to ascertain its limitations. In Figure 11, we find that our model is unable to detect Gazprom spokesman Sergei Kupriyanov, who is not a very prominent person. This follows a general pattern that entities on the long tail are, at times, wrongly predicted. Further, we see in Figure 15 that the model wrongly mistakes UN Conference for Chinese Embassy. This is another common mistake wherein the model fails to detect entities that have non-specific visual details in the video.

## G Further Analysis of Detected Entities

To detect named-entities from video, we train an Entity Perceiver (EP). In our implementation, EP is a GIT model trained on ground-truth entities extracted from video bullet summary using a LLM. We analysed the performance of EP in table 6 and fig. 7. In this section, we compare EP against state-of-the-art LLM based Vision-Language models (VLMs): BLIP-2 (Li et al., 2023a) and PaLI-X

(Chen et al., 2023). These models have demonstrated strong zero-shot results on a number of vision-language benchmarks. As such, we use them to detect entities and compare how our proposed EP fares against them.

To do so, we first sample representative frames from the video. Next, we input them to BLIP-2 and PaLI-X and prompt it predict the salient event, person, location and depicted organization in the video. A union of predictions from all the frames is collected as the final prediction. We compare the results against EP in fig. 16. Our initial exploration with these VLMs returned poor results. As such, we just demonstrate the performance comparison with our EP qualitatively. BLIP-2 is bullish in its predictions, wrongly detecting person Saddam Hussein (row 1) and Aamir Khan (row 4), and getting location and organization wrong in almost all samples. On the other hand, PaLI-X makes much less predictions and in the process doesn't predict person or organization in any sample. In contrast, EP correctly detects a number of named entities (Yasser Arafat, row 1; South Korea, US and Korean marnies, row 3; Bangladesh, row 4) while making much lesser false predictions. This highlights the superior performance of our trained EP in comparison to existing SOTA Vision-Language models.

	Event: Meeting Person: Terje Roed-Larsen, Yasser Arafat Organization: United Nations Location: Middle East, Palestinian GT	Event: Iraqi President meets Prime minister Person: Saddam Hussein, Osama Bin Laden Organization: Iraqi Government, Iraqi Army Location: Iraq BLIP
	Event: Meeting, Discussion Person: Yasser Arafat Organization: Israeli, Palestinian Location: Ramallah Pred	Event: Meeting, Protest Person: Organization: Location: Middle East Pali
	Event: Train Crash, Derailment, Visit Person: Spanish Prime Minister Mariano Rajoy Organization: Location: Site, Bridge GT	Event: Train Derails Location: Philippines Organization: San Francisco - San Bernardino Person: San Francisco, California BLIP
	Event: Train Derailment Person: Organization: Train Location: Portugal Pred	Event: Train Crash Person: Organization: Location: USA Pali
	Event: Joint military exercise, Landing Person: Organization: Korea Marines, US Marines Location: Pohang, Pyeongtaek GT	Event: War, Destroys Person: Organization: Iraqi Military, US Armed Forces Location: Iran, South Korea BLIP
	Event: Joint military exercise Person: Organization: Korea Marines, US Marines Location: South Korea Pred	Event: Amphibious assault, War, Landing Person: Organization: Location: US, Middle East Pali
	Event: lightning, wedding, death, injury, mourning Person: Mominul Haque Organization: Chapainawabganj District Hospital Location: Bangladesh GT	Event: Coronavirus, Treatment Person: Aamir Khan, Aafia Khan Organization: Pakistan Location: India BLIP
	Event: Killing, Injuring, Earthquake Person: Three people Organization: Police Location: Dhaka, Bangladesh Pred	Event: Ambulance, Funeral, Surgery, Covid Person: Organization: Location: US, Bangladesh, India Pali

Figure 16: Qualitative comparison of our Entity Detection module against BLIP (BLIP-2) and Pali (Pali-X).




Figure 17: Impact of entities on extracted contextual knowledge.

Additionally, we also quantitatively analyse GIT’s performance as Entity Perceiver. To this end, we compare GIT’s output to a reference baseline constructed using CLIP (Radford et al., 2021). To build this reference baseline, we query video-frames to a pretrained CLIP to retrieve the closest entity from train-set entities. Entities retrieved for all the frames are joined together to form the final predictions. We sample 6 frames for each video for

Model	Ent. F <sub>1</sub>	B-4	R-L
CLIP	0.02	1.14	16.30
GIT	<b>11.71</b>	<b>27.22</b>	<b>51.86</b>

Table 13: Quantitative comparison of detected entities against reference baseline – CLIP. GIT outperforms CLIP by a fair margin.

this baseline. We report the results on validation set in Table 13. From the table, we see that GIT outputs entities that are clearly better than CLIP. The result provides further evidence of high quality entities detected by our EP (GIT).

We also analyse how entities quality impacts contextual knowledge and in fig 17. The figure shows that quality of context improves with the entities generally, although it plateaus towards the end. This implies that we may need to make a lot of improvement in entity detection to see a significant improvement in the quality of context.

**Hallucination Study.** While Ent. F<sub>1</sub> provides some measure of hallucinated entities in the gener-

Model	Human ↓	GPT-4 ↓
GIT	64	79.2
+ $\mathcal{V}\mathcal{I}$ (Ours)	<b>60</b>	<b>78.6</b>

Table 14: Entity Hallucination Study. Lower is better. Our model consistently hallucinates less in both human and GPT-4 evaluation.

ated caption, we seek to explicitly quantify entities’ hallucination to better understand the model’s behavior. To this end, we measure hallucinated in predicted captions via a manual evaluation and a GPT-4 (Team, 2024) evaluation.

For human evaluation, two expert annotators not associated with this work are asked to detect if the caption predicted by the model contains hallucinated entities. To detect hallucination, we use strict criteria that even if there is a single entity in the predicted caption that is not present in the ground-truth caption, the predicted caption should be marked as hallucinated. A total of 100 random samples are evaluated. The results are reported in Table 14. We find that our model hallucinates less as compared to the baseline. Further, the strict hallucination criteria leads to a relatively high hallucination percentage for both the baseline and our model.

For GPT-4 evaluation, we input the ground-truth caption and the predicted caption to GPT-4 and task to detect it to detect hallucination in the predicted caption. Again, we use the strict criteria that even a single entity in the prediction caption that is not present in the ground-truth would make it hallucinated. Even for this evaluation, we find from Table 14 that our model performs better than the baseline. A total of 500 random samples were used for this evaluation.

## H Generalization across Dataset and Modality

**Generalization across Dataset** As mentioned before, existing news video captioning datasets are too small and the captions are not well aligned with the video content. However, to test the limits of our approach, we also report results on one of them – MM-AVS. Other datasets are not publicly available (Tang et al., 2022), not accessible (Li et al., 2020), or not in English (Krubiński and Pecina, 2023; Li et al., 2020). MM-AVS contains data from CNN and Daily Mail and has a total of 2173 samples. We report our results in table 15. As can be seen, our approach leveraging context outperforms vanilla

Model	BLEU-4	Rouge-L	CIDEr
M <sup>2</sup> SM (Fu et al., 2021)	-	30.23	-
SCCS (Qiu et al., 2022)	-	34.21	-
GIT	7.42	16.34	51.53
+ $\mathcal{V}\mathcal{I}$ (Ours)	<b>7.91</b>	<b>16.55</b>	<b>55.83</b>

Table 15: Evaluation on MM-AVS, a video news captioning dataset. Greyed rows use paired news article as input also.

Model	Ent. F <sub>1</sub>	BLEU-4	Rouge-L	CIDEr
Tough-to-Beat (Biten et al., 2019)	4.9	1.7	13.2	12.4
Pooled Embeddings (Biten et al., 2019)	5.3	2.1	13.5	13.2
Visual News (Liu et al., 2021)	18.59	5.3	17.9	50.5
Tell (Zhou et al., 2022)	23.44	11.6	25.0	107.6
GIT	6.9	5.19	15.63	26.82
+ $\mathcal{V}\mathcal{I}$ (Ours)	6.9	<b>5.24</b>	<b>15.78</b>	<b>27.93</b>

Table 16: Evaluation on Visual News, an image news captioning dataset. Greyed rows use paired news article as input also.

video only captioning by 4 CIDEr points. This illustrates the generalization of our approach to other news video datasets.

**Generalization across Modality.** We seek to understand if our approach can generalize to News Image Captioning as well. As such, we evaluate our approach on the largest news image captioning dataset: Visual News. We see table 16 that our approach improves our image-only training in this image domain as well. In fact, our approach using context comes close to the performance of some existing approaches that additionally use paired articles (greyed rows) to generate entity-aware captions.

## I Case Study on Human Performance on the Dataset

We present a challenging task in this work: generating entity-aware captions directly from the videos. To validate that this task is indeed practical, we do a case study where we ask human annotators to write entity-aware captions given videos. The annotators are allowed to use any external source: web, Wikipedia, LLMs etc. The annotators are technical people not associated with this work. The case study is visualized in Figure 18.

From this case study, we make the following observations:

- In general, humans can write good entity-aware captions using external sources. Hence, the task is practical.





	<p>At a press briefing, the US Defense Secretary <b>Donald Rumsfeld</b> and other officials discussed the situation in <b>Iraq</b>. Rumsfeld stated that while there are difficulties, the Iraqi people are better off now than they were before. The US civilian administrator of Iraq, <b>Paul Bremer</b>, emphasized that they are working to establish a sovereign and independent Iraqi government.</p>	<p><b>Donald Rumsfeld</b> and <b>Paul Bremer</b> giving a speech about the current state of <b>Iraq</b> at a press conference of the <b>Coalition Provisional Authority</b>. Paul Bremer states that they are working on restoring the sovereign and independence of the government. Donald Rumsfeld states the CPA is not responsible for protecting the Iraqi people but the Iraqi people should step up to protect themselves.</p>
	<p>A car bomb exploded outside an <b>Olympic bid arena</b> after a warning from ETA. The explosion caused smoke and shook the floor of nearby buildings. The area was cordoned off by police and fire engines.</p>	<p>An explosion occurred in <b>Madrid, Spain</b>, near the <b>Wanda Metropolitan Stadium</b>. Fire department have secured the scene. Locals say the explosion can be heard but no visible flames, only smoke.</p>
	<p>A senior Israeli defense official recommended that the evacuation of an outpost be delayed. Prime Minister <b>Ariel Sharon</b> convened a meeting with other ministers to discuss the issue, and stated that the evacuation of unauthorized outposts was approved by the <b>Israeli</b> government.</p>	<p>Israeli prime minister <b>Ariel Sharon</b> joins a meeting at the The Prime Minister's Office of <b>Israel</b>. The meeting discusses the Gaza Disengagement Plan for the Israeli army and settlement to retreat from <b>Gaza</b>.</p>
	<p>The morning after parliament elected a new president, scenes of Beirut showed a large billboard of the newly elected president, <b>Michel Suleiman</b>, on the side of a high-rise building... A Beirut resident expressed hope that things would return to normal and that <b>Lebanese</b> people who had left the country would want to come back and visit...</p>	<p>New <b>Lebanese</b> president <b>Michel Suleiman</b> was recently elected. His policies focused on rebuild and unification of the country after war.</p>
	<p><b>Ben Bernanke</b>, the <b>Federal Reserve</b> Chairman, discussed the current economic environment with the Senate Banking Committee. Bernanke signaled continued support for low interest rates, and noted that monetary policy is providing important support to the recovery while keeping inflation close to the FOMC's percent objective.</p>	<p><b>Ben S Bernanke</b>, Chair of the <b>Federal Reserve of the United States</b> talking about how US handling various economic problems.</p>
	<p>The inaugural session of the <b>summit</b> was attended by various heads of state. The Bhutanese Prime Minister opened the proceedings, while the Pakistani Prime Minister spoke about the need for cooperation to eliminate terrorism. The Afghan President discussed his country's recent strategic partnership agreement with India...</p>	<p>Video is about <b>Saarc summit 17</b> which focuses on building bridges. Representatives of multiple countries talk about how their countries are achieving this and plan to achieve it.</p>

Figure 18: Case study on random video samples to determine the practicality of the task. In general, humans are able to write reasonable entity-aware captions.

- Humans are able to detect all entities from ground-truth.
  - The context of the annotated captions are quite good and aligns strongly with the ground-truth caption.
  - At times, humans can detect entities that are not even mentioned in the ground-truth. For example Coalition Provisional Authority and Wanda Metropolitan Stadium in row 1 and row 2 respectively.
- From these observations, it's safe to conclude that while the task is challenging, it's practical.