Video2Commonsense: Generating Commonsense Descriptions to Enrich Video Captioning

Zhiyuan Fang* Tejas Gokhale* Pratyay Banerjee Chitta Baral Yezhou Yang Arizona State University, Tempe AZ

> {zy.fang, tgokhale, pbanerj6, chitta, yz.yang}@asu.edu https://asu-active-perception-group.github.io/Video2Commonsense

Abstract

Captioning is a crucial and challenging task for video understanding. In videos that involve active agents such as humans, the agent's actions can bring about myriad changes in the scene. Observable changes such as movements, manipulations, and transformations of the objects in the scene, are reflected in conventional video captioning. Unlike images, actions in videos are also inherently linked to social aspects such as intentions (why the action is taking place), effects (what changes due to the action), and attributes that describe the agent. Thus for video understanding, such as when captioning videos or when answering questions about videos, one must have an understanding of these commonsense aspects. We present the first work on generating commonsense captions directly from videos, to describe latent aspects such as intentions, effects, and attributes. We present a new dataset "Video-to-Commonsense (V2C)" that contains $\sim 9k$ videos of human agents performing various actions, annotated with 3 types of commonsense descriptions. Additionally we explore the use of open-ended video-based commonsense question answering (V2C-QA) as a way to enrich our captions. Both the generation task and the QA task can be used to enrich video captions.

1 Introduction

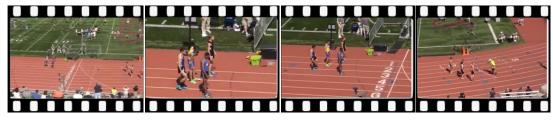
When humans watch videos they can typically understand and reason about various aspects of the scene beyond the visible objects and actions. This involves understanding that some objects are *active agents* that not only perform actions and manipulate objects, but are motivated by intentions, have pre-conditions, and that their actions have an effect on the world and their own mental states. For instance, in analyzing the video clip in Figure 1,

humans employ various capabilities such as perception, reasoning, inference, and speculation, to come up with a description for the observable sequence of events, but also reason about latent aspects such as the intention of the group of runners "to win the medal", the effect of being "congratulated at the finish line", and the attribute "athletic".

The above example also illustrates that recognition of objects, actions, and events is often not enough; understanding causal relationships, social interactions, and commonsense aspects behind them provides context and a more semantic interpretation of the video (Gupta et al., 2009). A model that can provide such detailed interpretations facilitates answering inferential questions, such as "Will the player get angry later?". However, existing visual understanding systems are unable to perform such tasks that require speculative reasoning. A critical missing element in complex video understanding is the capability of performing commonsense inference, especially a generative model. Existing efforts seek to find textual explanations or intentions of human activities as a classification task (Vondrick et al., 2016) or a vision-to-text alignment problem (Zhu et al., 2015).

In this paper we propose the Video to Commonsense (V2C) framework to generate visually grounded commonsense descriptions about the underlying event in the video, enriching the factual description provided by a caption. Under this framework a system is expected to generate captions as well as three types of commonsense descriptions (intention, effect, attribute) directly from an input video. The V2C model can also be used as a building block for downstream tasks such as video question answering for questions requiring commonsense. Inspired by (Bosselut et al., 2019), our model – the "V2C-Transformer" utilizes: (1) a video encoder to extract global representations of the video, (2) a transformer decoder that generates

^{*}Equal Contribution



Conventional Caption

Group of runners get prepared to run a race.

Commonsense-Enriched Caption

In order to win a medal, a group of runners get prepared to run a race. As a result they are congratulated at the finish line. They are athletic.

Commonsense Question Answering

What happens next to the runners?

Are congratulated at the finish line become tired

Figure 1: Comparison of conventional video captioning with our commonsense-enriched captioning. Our captions describe intention behind the action (red), attribute of the agent (blue), and effect of the action on the agent (green).

captions and commonsense descriptions, and (3) a cross-modal self-attention module that exploits joint visual-textual embeddings.

We curate the V2C dataset for training and benchmarking models on this task. We adopt the MSR-VTT video description dataset (Xu et al., 2016) as a source of videos and captions. We first utilize the ATOMIC machine commonsense dataset (Sap et al., 2018) to get a list of candidate commonsense texts (intentions, effects, and attributes), and rank these using a BERT-based (Devlin et al., 2019) model. Since these candidates are retrieved without using the video and may not be accurate, we instruct humans to watch the videos and select, remove, or rewrite the texts retrieved from ATOMIC. The text retrieved by ATOMIC helps our human annotators to understand the format of desired annotations, and also gives them a list of suggestions. The human component in our annotation procedure makes our data visually grounded and relevant, linguistically diverse, and natural.

We additionally explore the use of our V2C-Transformer architecture for a open-ended video question answering task, where the questions are about commonsense aspects from the video. For this, we create a QA addendum of the V2C dataset called V2C-QA. By asking questions about the latent aspects in the video, our models are able to enrich caption generation with three specific types of commonsense knowledge.

Our contributions are summarized below:

- 1. We formulate the "V2C" task for enriching video captioning by generating descriptions of commonsense aspects.
- 2. We curate a video dataset annotated with captions and commonsense descriptions.

- 3. We present our V2C-Transformer architecture that generates relevant commonsense descriptions, and serves as a strong baseline.
- 4. We pose V2C as a video question answering task and show that it can assist commonsense caption generation.

2 Video to Commonsense (V2C)

Problem Formulation: Consider a video V consisting of N_v frames described by sentence S. Our Video-to-Commonsense (V2C) framework can be used for generating commonsense descriptions C under two settings. In the first setting (V2C-Completion), we use ground-truth captions to guide commonsense-enriched caption generation. This task can be viewed as providing supplementary explanations to the caption. In the second setting (V2C-Generation), we first learn to generate captions from videos, $\mathbf{g}(V)$, and then use them to generate commonsense descriptions.

V2C-Completion
$$C = \mathbf{f}(V, S)$$
.
V2C-Generation $C = \mathbf{f}(V, \mathbf{g}(V))$. (1)

2.1 V2C-Transformer

The proposed Video2Commonsense Transformer is a cross-modal model that generates captions and commonsense-enriched descriptions from videos. Our approach (Figure 2) adopts the "encoder-decoder" design: a video encoder that extracts global representations of the input video, and a transformer decoder that produces relevant commonsense knowledge along with captions.

Video Encoder: We obtain per-frame ResNet-152 (He et al., 2016) features for video V and process them using an LSTM model (Sundermeyer

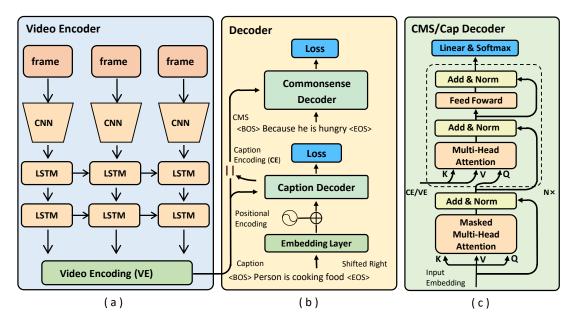


Figure 2: The V2C-Transformer model architecture contains: (a) Video Encoder designed to take video frames as input and encode them into frame-wise representations, (b) Decoder module consisting of a Caption Decoder and a Commonsense Decoder, and (c) Transformer Decoder module containing a stack of N consecutive transformer blocks (shown inside the dashed area).

et al., 2012), a standard architecture for modeling long temporal sequences, and use the last hidden states of the LSTM as the video representations. We concatenate all previous hidden states from each LSTM module as a final global video encoding v, to provide the model with explicit context using the temporal attention mechanism.

Decoder: The video encoding is used as input to two decoder networks that use a transformer language model (Radford et al., 2018) to generate a caption and commonsense description, using an inference mechanism similar to Bosselut et al. (2019). Our model is a two-stage process that first predicts the current events directly from videos, and then produces the corresponding commonsense captions. During training, the caption decoder \mathbf{D}_{CAP} takes the video encoding (v) and ground truth caption (s) as input to generate caption encoding (ŝ), while the commonsense decoder \mathbf{D}_{CMS} uses the concatenation of video and caption encoding to obtain the commonsense description (c), as shown in Figure 1 (b). This arrangement enables the attention module in commonsense decoder to attend to both the video and caption context.

$$\mathbf{\hat{s}} = \mathbf{D}_{CAP}(\mathbf{v}, \mathbf{s}), \quad \mathbf{c} = \mathbf{D}_{CMS}(\mathbf{v}, \mathbf{\hat{s}}). \tag{2}$$

Transformer Decoder is composed of a stack of transformer blocks (dashed area in (c) Figure 2),

whose main component is a self-attention architecture. It takes as input the summation of word embedding and the positional encoding offset by 1 position through masked multi-head attention, which prevents the future words been seen. In our model, we deploy two stacked decoder architectures for both caption decoding and commonsense knowledge decoding. The Transformer Block consists of consecutive linear transformation: a multi-head attention module (denoted as $\mathcal{H}_{\text{M-ATT}}$), a two-layer feed forward network (\mathcal{H}_{FFN}), a layer normalization operation, and a residual connection.

Multi-head Attention module To enable our transformer decoder to generate commonsense descriptions by using both the visual and textual content, we modify the multi-head attention module (which acts as the basic unit in recent transformer based language generation models (Radford et al., 2018, 2019)) as a cross-modal module. \mathcal{H}_{M-ATT} takes the input of the embedding of key (K), value (V) and query (Q). The key and value in transformer block are the video encoding (caption decoder) or concatenation of video/caption encoding (commonsense decoder), while the query is the output from the previous transformer block. In the masked multi-head attention module, K, V and Q are the identical vectors of input embedding. For a self-attention block with h heads,

$$\mathcal{H}_{M-ATT}(K, V, Q) = \mathcal{H}_{FFN}([x_1, \dots, x_h]), \qquad (3)$$

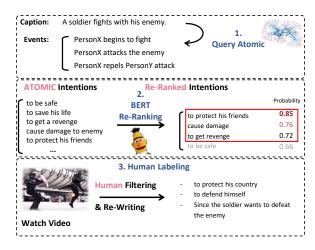


Figure 3: The overall three-step pipeline (retrieval from ATOMIC, BERT re-ranking, and human labeling) to construct our V2C dataset.

where x_i is computed by scaled dot-product attention operation, for head-index i, key-dimension d_k n, and transformation parameters W_i .

$$\begin{split} & \textbf{for } \mathbf{D}_{\text{CAP}}, \quad x_i = \text{Softmax}\big(\frac{\mathbf{W}_i^{\text{Q}} \mathbf{Q} \cdot \mathbf{W}_i^{\text{K}} \mathbf{K}'}{\sqrt{d_k}}\big) \mathbf{W}_i^{\text{V}} \mathbf{V}, \\ & \textbf{for } \mathbf{D}_{\text{CMS}}, \quad x_i = \text{Softmax}\big(\frac{\mathbf{W}_i^{\text{Q}} [\mathbf{v}, \mathbf{s}] \cdot \mathbf{W}_i^{\text{K}} [\mathbf{v}, \mathbf{s}]'}{\sqrt{d_k}}\big) \mathbf{W}_i^{\text{V}} \mathbf{V}. \end{split}$$

3 The V2C Dataset

For the V2C task we need video clips annotated with commonsense descriptions about the agents in the video, as shown in Figure 1. While there are video captioning datasets such as MSR-VTT (Xu et al., 2016), the captions in these datasets describe only the observable objects in the image, but do not describe latent and commonsense aspects. We are the first to curate such a dataset with annotations describing the intention of agent to perform an action, the effect of the action and the attribute of the agent given the action.

MSR-VTT contains around 10k videos each 10 to 30 seconds long, belonging to 20 categories covering a variety of topics such as sports, music, news, and home videos. Each video is accompanied by 20 human-annotated textual descriptions on average. For training and benchmarking the novel V2C task, we further complement MSR-VTT with event-level commonsense annotations, i.e. event descriptions with intentions, effects and attributes. We remove captions and videos that do not have clear human activities. This is because having such videos leads to an imbalance in the number of captions for each

Type	Video Caption	Commonsense
Intention	Two guys are wrestling A man and woman are singing	to beat the opponent to express them- selves musically
Attribute	A guy is singing in a crowd Group of riders race on motorcycles.	outgoing adventurous
Effect	A person is making a paper airplane A man and a woman are talking to each other	gets excited to fly it share ideas and opin- ions

Table 1: Examples of commonsense annotations (intentions, attributes and effects) retrieved from ATOMIC for captions in MSR-VTT.

video, thus making it inappropriate to just evaluate caption generation using BLEU scores.

ATOMIC (Sap et al., 2018) is an atlas of every-day commonsense knowledge and contains 880k triplets about causes and effects of human activities, organized as *if-then* relations, annotated by crowd-sourced workers. This data can be categorized based on causal relations, thereby giving us the categories "cause", "effect" and "attribute", e.g., "*if* X wants to relax, *then* he will play video game."

3.1 Querying from ATOMIC and Re-ranking

Since inferential knowledge in ATOMIC only covers human activities, we first retain only those captions in Msr-vtt that describe human activities. We then select three queries from ATOMIC most similar to the caption, and extract the commonsense descriptions corresponding to these queries. In order to select a more reasonable subset of commonsense descriptions, we first train a ranking model. We use the BERT (Devlin et al., 2019) architecture for the ranking model, trained on the ATOMIC dataset for a binary classification task, to predict the relevance of a candidate commonsense description with respect to the event. We select the top three relevant intentions, effects, and attributes for each caption. This allows us to obtain a preliminary set of 9 commonsense annotations per video directly from the ATOMIC dataset, relevant to the caption, albeit with noise and annotations that are not relevant to the video.

3.2 Detailed Human Annotation

Since we do not use the video to retrieve commonsense descriptions from ATOMIC, we employ human workers to annotate our dataset. We recruit

Relation	Model	CIDER	$\mathbf{PPL}\downarrow$	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR	ROUGE-L
	S2VT (Venugopalan et al., 2015)	-	-	35.9	-	-	-	-	-
A 44	Attention-Enc-Dec (Gao et al., 2017)	-	-	38.3	-	-	-	-	-
Attribute	Dense Captioner (Zhou et al., 2018)	-	-	46.0	-	-	-	-	-
	Video CMS Transformer	-	-	47.3	-	-	-	-	-
	S2VT (Venugopalan et al., 2015)	28.3	23.6	24.9	18.6	16.2	14.3	15.4	22.1
Effect	Attention-Enc-Dec (Gao et al., 2017)	29.5	22.0	26.5	19.4	18.8	15.1	17.5	23.9
Effect	Dense Captioner (Zhou et al., 2018)	36.9	16.0	33.7	24.8	21.0	20.2	20.0	29.9
	Video CMS Transformer	37.3	15.6	34.8	25.9	22.5	20.4	20.8	30.6
	S2VT (Venugopalan et al., 2015)	51.8	17.8	48.4	39.9	34.3	26.4	23.3	44.3
Intention	Attention-Enc-Dec (Gao et al., 2017)	52.1	16.0	51.1	42.6	35.5	28.2	24.3	48.0
	Dense Captioner (Zhou et al., 2018)	60.3	12.0	59.3	47.0	37.3	31.5	28.0	53.1
	Video CMS Transformer	62.0	11.7	60.8	48.4	39.1	34.1	28.5	54.6

Table 2: Evaluation of V2C completion task using CIDER, BLEU, Perplexity, Rouge, and Meteor metrics. We use only BLEU-1 to evaluate the attribute generation since the average length of the ground truth is just less than 2.

two sets of human workers to watch the video, read the caption and select/annotate the relevant commonsense descriptions for each video. The first set is Amazon Mechanical Turkers (AMT) who select relevant descriptions. The second set is skilled human annotators, screened from a set of university students proficient in English, who are asked to provide annotations in their own words, and remove or edit irrelevant annotations that were provided by ATOMIC and AMT workers. This makes our annotations not only grounded in the video, but also more descriptive, linguistically diverse, and of higher quality (see Figure 3). The descriptions from ATOMIC, although not relevant to the video in some cases, give our workers an idea about the format of annotations desired. The skilled humans reported that 95% of the captions were relevant, and 65% of the ATOMIC descriptions were useful in understanding the annotation task. Through this procedure, we obtain 6819 videos for training and 2906 videos for testing, a total of 121,651 captions $(\sim 12 \text{ captions/video})$, each caption accompanied with 5 commonsense knowledge annotations (V2C-Raw set). In experiment, we use video captioning technique to conduct the V2C completion task on V2C-Raw set. In addition, we instruct human annotators to select and rewrite one raw phrase into complete sentences that complement the captions. In total we have 3 complete sentences per video for intention/effect/attribute respectively, and this yields a subset that allows our model to generate complete story-like sentences (V2C-Clean Set). Table 1 shows examples from the newly compiled dataset. We conduct rigorous human evaluation to evaluate the quality of our V2C dataset ("Gold Annotations" in Table 3). Details about the dataset creation process and quality control mechanisms can be found in the Appendix.

4 Experiments

In this section we describe the loss function used for training our model, additional details about video pre-processing, hyper-parameters, and baseline models, and the metrics used for evaluation.

Loss Function: The decoder parameters Θ are trained to maximize the log-likelihood over the training set given by $\mathcal{L} = \mathcal{L}_{cap} + \mathcal{L}_{cms}$, where

$$\begin{split} \mathcal{L}_{cap} &= \sum_{t=1}^{N_S} \log \Pr(\mathbf{y}_t|\mathbf{y}_{t-1}, \mathbf{v}; \mathbf{\Theta}), and \\ \mathcal{L}_{cms} &= \sum_{t=1}^{N_C} \log \Pr(\mathbf{y}_t|\mathbf{y}_{t-1}, [\mathbf{v}, \widetilde{\mathbf{s}}]; \mathbf{\Theta}). \end{split} \tag{4}$$

 \mathbf{y}_t denotes the one-hot vector probability of each word at time t, and N_S, N_C denote the length of the caption and commonsense respectively.

Setting: In order to obtain video representations, we uniformly sample 40 frames from each video and extract features using feed ResNet (He et al., 2016) pre-trained on Imagenet ILSVRC12 dataset (Deng et al., 2009) and get a 2048-d output from the last layer. We use one-hot input (1-of-*N* encoding) of the text input and pass it through an embedding layer to produce a 1028-d hidden vector. We use independent vocabularies for captioning and commonsense generation with sizes 27,603 and 24,010 respectively. Note that, as the generated

Hyperparameters: Our decoder is a lightweight transformer decoder consisting of 6 transformer blocks with 8 attention heads each. We use Adam optimizer with 5000 warm-up steps, and learning rate initialized at 1e-4, and a dropout probability of 0.1 after the residual layer. Our model is trained on a machine with single NVIDIA 1080-Ti GPU.

Task	Task Model		Attribute	Intention	Average	Caption
E2C-Completion	9ENC9DEC (Sap et al., 2018)	44.23	52.01	49.72	49.47	-
(Text-Only)	COMET (Bosselut et al., 2019)	54.98	56.28	66.32	59.22	-
V2C-Completion	Att-Enc-Dec(Gao et al., 2017)	66.09	52.40	56.26	58.25	-
v2C-Completion	VCT-Completion	<u>66.83</u>	63.45	<u>67.37</u>	65.88	-
V2C-Generation	Att-Enc-Dec(Gao et al., 2017)	55.93	74.87	65.54	64.78	74.67
v 2C-Generation	VCT-Generation	62.99	73.54	66.74	<u>67.76</u>	73.17
Gold Annotations	V2C Dataset	75.19	83.03	80.11	79.44	95.01

Table 3: Human evaluation scores for V2C. Captions are an input for the **V2C-Completion** task, and generated for the **V2C-Generation** task. The best model is given in bold, while the overall best is underlined.

Baseline Model: We compare our method with strong video captioning baseline models like, S2VT (Venugopalan et al., 2015), "Attention-Enc-Dec" (Gao et al., 2017) - LSTM based models which reach competitive performing on MSR-VTT dataset. and "Dense Captioning" (Zhou et al., 2018), which is a transformer based video captioning model. As "Dense Captioning" is proposed to generate multiple continuous captions for a long untrimmed videos, we modify this by removing the temporal bounding boxes prediction module, and produce two continuous captions (caption + commonsense sentence) together without corresponded starting and ending time. All baselines are trained to predict commonsense descriptions from video on the V2C dataset. We do not compare with VideoBERT (Sun et al., 2019) which is trained on a limited set of cooking videos and hence nontransferable, and requires individual captions for multiple segments of the video.

Metrics: We report both the performances evaluated by automatic scores and human evaluations following the protocols from (Bosselut et al., 2019; Sap et al., 2018). We evaluate our method using BLEU (n=1-4) (Papineni et al., 2002), Meteor (Banerjee and Lavie, 2005), Rouge (Lin, 2004), and perplexity score of the generation on its corpus. We further conduct human evaluations using AMT workers, who are asked to identity whether the generated commonsense justifiably completes the events (V2C-completion). We follow the setup in (Sap et al., 2018) and randomly sample 100 videos from test set and collect 10 generations for each. To guarantee the objectiveness of the human evaluations, we hire 5 workers for each sample, yielding 30k ratings in total for each model.

4.1 Results

Natural Language Generation Metrics: We show evaluation of the commonsense comple-

tion task in Table 2. Compared to the baseline model, our method exhibits a consistent and overall improvement on almost all metrics. Our V2C-Transformer significantly outperforms the LSTM based model in (Gao et al., 2017) by 7.7% at BLEU-4 for the intention prediction. Because the V2C-Transformer and the LSTM model share a similar video encoder, our performance improvement could be attributed to the use of self-attention mechanisms in the transformer block in decoding phase. This observation is consistent with the conclusion from (Bosselut et al., 2019), and yields further support to the transformer architecture being suited for commonsense inference tasks. Moreover, when compared with DenseCap which has a similar transformer architecture and parameters, our model exhibits better evaluation scores, verifying it as a strong baseline model for the V2C task.

Human Evaluation In Table 3, E2C (Event to Commonsense) is the task of commonsense completion given only textual events (Sap et al., 2018; Bosselut et al., 2019) as opposed to V2C which uses both text and video. 9ENC9DEC (Sap et al., 2018) is composed of nine GRU based encoder-decoders as a baseline model for commonsense completion on text, and COMET (Bosselut et al., 2019) is a large-scale generative pre-trained transformer (GPT) model (Radford et al., 2018). We would like to highlight that our transformer model is light-weight with only half of the parameters in GPT without any pre-training.

We evaluate our model on the tasks of caption generation with human evaluations, and also compare it with the gold annotations. Our gold annotation for ground-truth captions (sourced from the MSR-VTT dataset) points to the fact that a small percentage of captions from MSR-VTT are not relevant to the video, and this is amended by our human workers.

For the V2C-Completion task, our V2C-

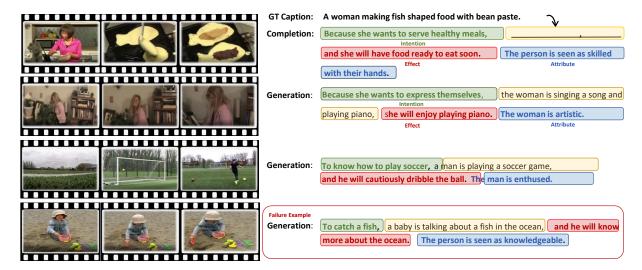


Figure 4: Examples of outputs of our model for the V2C Completion and Generation tasks along with the ground-truth (GT) caption. A failure example shown in the bottom red box.

Transformer model is substantially better (by 7.73%) than the LSTM-based model from (Gao et al., 2017), and shows consistent lead on each dimension. Thus, when the ground-truth caption is given, our model is able to generate much more relevant commonsense descriptions, thereby consolidating it's ability of commonsense generation.

For the task of V2C-Generation, the difference between human scores for LSTM vs V2C-Transformer is reduced, but our VTC outperforms on average by 2.98%. This may be attributed to the fact that the LSTM-based model is slightly better at generating captions.

Generating Textual Stories with Commonsense

In order to generate story-like textual descriptions that complement the factual captions, we additionally train our model to exploit our diverse completesentence annotations. Specifically, instead of producing the commonsense knowledge given the videos and captions, we finetune our pre-trained V2C-Transformer model on predicting the human rewritten texts, and generate complete story-like captions. Since we do not have enough annotations per sample to compute a fair BLEU score for comparisons, we showcase some sample generated descriptions for qualitative analysis (see Figure 4). With that, we observe V2C-Transformer is able to produce complete stories that contain simple, while logically consistent storylines that complement both the visual content and the factual descriptions. We believe that collecting a set of story-like sentences will further enrich our models, and allow us to generate much more contextual, creative, and

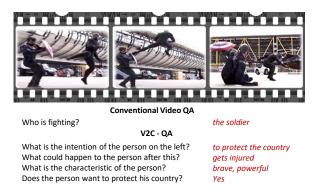


Figure 5: Example questions from V2C-QA compared with conventional video question answering.

natural commonsense descriptions from a video.

5 V2C-QA

Another way of generating commonsense descriptions about the video is by asking pointed questions. Consider the example in 1 where we ask the question "What happens next to the runners", about the effect of the action "prepare" performed by the agents "group of runners" observed in the video. We propose a V2C-QA – an open-ended commonsense video question-answering task, where we ask questions about the intents, effects and attributes of the agents in the video.

Dataset: We use the caption and commonsense annotations in the V2C dataset to create question-answer pairs for each video. We first extract the action and subject from the caption using SpaCy linguistic features (Honnibal and Johnson, 2015). For each intention, attribute and effect for a video, we use template-based generation to get 7 types

	Model	top	-1	top-3		top-5	
	1710del	p	r	p	r	p	r
Ξ	MSR-VTT QA	9.68	2.13	7.15	4.68	6.07	6.60
į	V2C-T	10.34	2.31	7.69	5.03	6.37	6.87
Intention	V2C-T + Captions	10.72	2.54	8.08	5.47	6.39	7.20
王	Pretrained V2C-T	10.77	2.69	8.01	5.58	6.71	7.88
	Pretrained V2C-T + Cap.	11.04	2.68	7.96	5.70	6.63	7.79
	MSR-VTT QA	19.89	5.02	8.04	5.91	5.30	6.49
ಕ	V2C-T	20.95	5.43	8.65	6.57	5.65	7.06
Effect	V2C-T + Captions	20.95	5.32	8.50	6.48	5.76	7.26
Œ	Pretrained V2C-T	20.95	5.32	8.63	6.55	5.82	7.49
	Pretrained V2C-T + Cap.	21.12	5.60	8.70	6.89	5.83	7.68
	MSR-VTT QA	46.10	37.22	16.02	49.45	7.49	41.03
ūŧ	V2C-T	59.52	48.30	22.39	51.40	13.97	52.57
Ë	V2C-T + Captions	59.74	48.22	23.12	52.44	14.64	54.35
Αtt	Pretrained V2C-T	60.72	49.00	23.18	52.73	14.98	55.40
	Pretrained V2C-T +Cap.	59.57	48.24	23.10	52.54	14.94	54.91
	Text-Only Baseline	12.36	11.70	13.84	12.35	14.77	14.10

Table 4: Precision (p) and Recall (r) for V2C-QA for each type of question.

of questions – yielding 21 questions per sample, including negative questions as in Gokhale et al. (2020). In total, we have 1,250 training videos and 250 test videos, and a total of 37k questions. We have a set of 5,555 unique answers for our questions. Each question can have multiple possible true answers as shown in the example in Figure 5. The V2C-QA task asks questions that require commonsense reasoning about internal mental states, motivations, and latent aspects of agents in the video as opposed to the conventional video-QA questions about visible objects and actions.

Models: We utilize our V2C-Encoder followed by an open-ended answering module. We jointly predict the type of the question and combine it with the V2C encoding using a feed-forward network. For textual features, we use embeddings from BERT-base (Devlin et al., 2019). Our models are trained on the open-ended QA task and setup as a multi-label classification task similar to VQA (Antol et al., 2015), with an answering module design inspired by LXMERT (Tan and Bansal, 2019). Our loss function includes the classification loss for answering, the attention loss for question-type, and a label-ranking loss.

Results: MSR-VTT QA (Xu et al., 2017) is as a good baseline since it is trained on a conventional videoQA task on the MSR-VTT videos, and only takes video and query as input, unlike recent video understanding models (Lei et al., 2018) that take additional supervision, such as subtitles. However this model is trained for a multiple-choice QA

scheme, so we modify it with our open-ended answering module. We compare our models when we use our encoder pretrained on the V2C caption generation task, and then finetune it on the V2C-QA task. We also train models with ground-truth factual captions as input. Our results are shown in Table 4, where we evaluate on prediction of top-k (1,3,5) answers, and report precision and recall. Our encoder pre-trained on the V2C task outperforms all other models. Attribute-related questions are easier to answer, while the models struggle the most for questions about intention. Captions help in questions about effects. The overall text-only baseline shows an insignificant bias between the question and answer-options.

6 Related Work

Video Captioning: Captioning is crucial for understanding visuals; however it is typically limited to describing observable objects and events (Yang et al., 2011; Thomason et al., 2014; Gan et al., 2017)), or for generating paragraphs or multisentence captions about the image or video (Krause et al., 2017; Krishna et al., 2017). However, for detailed video understanding, one needs to obtain descriptions that go beyond observable visual entities and use background knowledge and commonsense to reason about objects and actions. Work for inferring motivations of human actions in static images by incorporating commonsense knowledge are reflected in Pirsiavash et al. (2014); Vondrick et al. (2016). Commonsense caption generation has been approached on abstract scenes and clipart images in Vedantam et al. (2015). We present the first generative model for commonsense video captioning.

Video Question Answering: Since caption generation can only describe observable events, recent work seeks to move closer to comprehension, by learning to answer complex questions about videos. However, the datasets used for Video QA (Yang et al., 2003; Xu et al., 2016; Zhu et al., 2017) focus only on directly evident visual concepts and construct the questions mostly about "where" and "what" aspects. Question answering on movie videos has been explored by Tapaswi et al. (2016) who collect questions about "why" and "how" aspects. Recently Lei et al. (2018); Zadeh et al. (2019) have propose video-based QA tasks with open-ended high-order questions that need multimodal understanding, social intelligence modeling,

and spatio-temporal reasoning. We introduce a novel open-ended video question answering task in this paper, where the questions are about three aspects of commonsense human behavior.

Visual Reasoning: Aspects of visual reasoning have been explored by Yatskar et al. (2016) as a situation recognition task on single images, and in Visual Madlibs (Yu et al., 2015) as a "fill-inthe-blanks" task for single-image captioning that contains some categories which require reasoning about internal mental states and future events. Kim et al. (2018) provide textual explanations for actions in a self-driving scene. Zellers et al. (2019) propose a visual question answering task that requires commonsense reasoning to answer a question and to provide a rationale behind the answer. Spatial and compositional reasoning is required to answer questions about synthetic images in CLEVR (Johnson et al., 2017). Critical aspects of visual reasoning also include the model's ability to conduct object grounding by natural language descriptions (Rohrbach et al., 2016; Fang et al., 2018, 2019). Another aspect of visual reasoning is the ability predict a sequence of actions (procedure planning), or to reason about intermediate video frames (walkthrough planning) between two frames, explored in Gokhale et al. (2019); Chang et al. (2020).

Textual Commonsense: Commonsense-based question answering is an area of active research with several datasets and challenges requiring reasoning about conceptual commonsense (Talmor et al., 2019), physical commonsense (Bisk et al., 2020), social commonsense (Sap et al., 2019), and abductive commonsense (Bhagavatula et al., 2020). On the other hand, challenges such as ProPara (Mishra et al., 2018) and bAbI (Weston et al., 2015) require tracking elements, actions, and effects of actions. Commonsense-based text generation has recently been explored via the ATOMIC dataset (Sap et al., 2018), a corpus of 877k textual descriptions of inferential knowledge organized as if-then relations. Bosselut et al. (2019) adopt the ATOMIC dataset to learn a generative model of commonsense knowledge. To the best of our knowledge, ours is the first work on generating commonsense descriptions from visual inputs.

7 Outlook

A video typically contains one or many objects (sometimes performing actions) in different back-

grounds, scenes, or situations. Some objects may be "passive" such as trees or buildings, while some objects may be "active" such as people performing actions like walking, singing, and driving. This paper is focused on describing such active agents in terms of their intentions, effects of their actions, and attributes that characterize these agents.

We distinguish V2C from the traditional video captioning task. Video captions describe observable objects, background, and actions, while commonsense descriptions in our task seek to describe the unobservable intentions of the agent (pre-conditions or mental conditions), effects of the action (that happen in the future), and attributes which characterize the agent. Thus commonsense generation goes *beyond the visible*. Ours is the first attempt at developing a generative video-based commonsense model. We anticipate that our framework can be utilized for many applications in video understanding, comprehension, human-robot interaction, and learning commonsense in a multimodal setting.

8 Conclusion

In this paper, we explore a novel and challenging task to generate video descriptions with rich commonsense descriptions that complement the factual captions. We expand an existing video captioning dataset for the V2C task through automated retrieval from a textual commonsense corpus followed by human labeling, and present a novel V2C-Transformer model to serve as a strong baseline method for the V2C task. Our evaluation verifies the effectiveness of our method, while also indicating a scope for further study, enhancement, and extensions in the future. Our experiments on using the V2C-Transformer as a component for the V2C-QA task show that the model has transfer learning capabilities that can be applied to other vision-andlanguage tasks such as question-answering, that require commonsense reasoning.

Acknowledgements

The authors acknowledge support from the NSF Robust Intelligence Program project #1816039, the DARPA KAIROS program (LESTAT project), the DARPA SAIL-ON program, and ONR award N00014-20-1-2332. ZF, TG, YY thank the organizers and the participants of the Telluride Neuromorphic Cognition Workshop, especially the Machine Common Sense (MCS) group.

References

- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433.
- Satanjeev Banerjee and Alon Lavie. 2005. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings* of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization, pages 65–72.
- Chandra Bhagavatula, Ronan Le Bras, Chaitanya Malaviya, Keisuke Sakaguchi, Ari Holtzman, Hannah Rashkin, Doug Downey, Wen-tau Yih, and Yejin Choi. 2020. Abductive commonsense reasoning. In *ICLR*.
- Yonatan Bisk, Rowan Zellers, Ronan LeBras, Jianfeng Gao, and Yejin Choi. 2020. Piqa: Reasoning about physical commonsense in natural language. In *AAAI*, pages 7432–7439.
- Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. 2019. Comet: Commonsense transformers for automatic knowledge graph construction. In *Association for Computational Linguistics*.
- Chien-Yi Chang, De-An Huang, Danfei Xu, Ehsan Adeli, Li Fei-Fei, and Juan Carlos Niebles. 2020. Procedure planning in instructional videos. *European Conference on Computer Vision*.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition, pages 248–255. IEEE.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In NAACL-HLT (1).
- Zhiyuan Fang, Shu Kong, Charless Fowlkes, and Yezhou Yang. 2019. Modularized textual grounding for counterfactual resilience. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6378–6388.
- Zhiyuan Fang, Shu Kong, Tianshu Yu, and Yezhou Yang. 2018. Weakly supervised attention learning for textual phrases grounding. *arXiv preprint* arXiv:1805.00545.
- Zhe Gan, Chuang Gan, Xiaodong He, Yunchen Pu, Kenneth Tran, Jianfeng Gao, Lawrence Carin, and Li Deng. 2017. Semantic compositional networks for visual captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5630–5639.

- Lianli Gao, Zhao Guo, Hanwang Zhang, Xing Xu, and Heng Tao Shen. 2017. Video captioning with attention-based lstm and semantic consistency. *IEEE Transactions on Multimedia*, 19(9):2045–2055.
- Tejas Gokhale, Pratyay Banerjee, Chitta Baral, and Yezhou Yang. 2020. Vqa-lol: Visual question answering under the lens of logic. In *European Conference on Computer Vision*.
- Tejas Gokhale, Shailaja Sampat, Zhiyuan Fang, Yezhou Yang, and Chitta Baral. 2019. Cooking with blocks: A recipe for visual reasoning on image-pairs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 5–8.
- Abhinav Gupta, Praveen Srinivasan, Jianbo Shi, and Larry S Davis. 2009. Understanding videos, constructing plots learning a visually grounded storyline model from annotated videos. In 2009 IEEE Conference on Computer Vision and Pattern Recognition, pages 2012–2019. IEEE.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on* computer vision and pattern recognition, pages 770– 778.
- Matthew Honnibal and Mark Johnson. 2015. An improved non-monotonic transition system for dependency parsing. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1373–1378, Lisbon, Portugal. Association for Computational Linguistics.
- Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing.
- Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. 2017. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 2901– 2910.
- Jinkyu Kim, Anna Rohrbach, Trevor Darrell, John Canny, and Zeynep Akata. 2018. Textual explanations for self-driving vehicles. In *Proceedings of the European conference on computer vision (ECCV)*, pages 563–578.
- Jonathan Krause, Justin Johnson, Ranjay Krishna, and Li Fei-Fei. 2017. A hierarchical approach for generating descriptive image paragraphs. In *Computer Vision and Patterm Recognition (CVPR)*.
- Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. 2017. Dense-captioning events in videos. In *International Conference on Computer Vision*.

- Jie Lei, Licheng Yu, Mohit Bansal, and Tamara L Berg. 2018. Tvqa: Localized, compositional video question answering. In EMNLP.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. arXiv preprint arXiv:1907.11692.
- Bhavana Dalvi Mishra, Lifu Huang, Niket Tandon, Wen-tau Yih, and Peter Clark. 2018. Tracking state changes in procedural text: A challenge dataset and models for process paragraph comprehension. In *Proceedings of NAACL-HLT*, pages 1595–1604.
- 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 on association for computational linguistics, pages 311–318. Association for Computational Linguistics.
- Hamed Pirsiavash, Carl Vondrick, and Antonio Torralba. 2014. Inferring the why in images. Massachusetts Inst of Tech Cambridge.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8).
- Anna Rohrbach, Marcus Rohrbach, Ronghang Hu, Trevor Darrell, and Bernt Schiele. 2016. Grounding of textual phrases in images by reconstruction. In *European Conference on Computer Vision*, pages 817–834. Springer.
- Maarten Sap, Ronan LeBras, Emily Allaway, Chandra Bhagavatula, Nicholas Lourie, Hannah Rashkin, Brendan Roof, Noah A Smith, and Yejin Choi. 2018. Atomic: An atlas of machine commonsense for ifthen reasoning. In *AAAI*.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. 2019. Social iqa: Commonsense reasoning about social interactions. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pages 4453–4463.
- Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. 2019. Videobert: A joint model for video and language representation learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7464–7473.

- Martin Sundermeyer, Ralf Schlüter, and Hermann Ney. 2012. Lstm neural networks for language modeling. In *Thirteenth annual conference of the international speech communication association*.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. Commonsenseqa: A question answering challenge targeting commonsense knowledge. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158.
- Hao Tan and Mohit Bansal. 2019. Lxmert: Learning cross-modality encoder representations from transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5103–5114.
- Makarand Tapaswi, Yukun Zhu, Rainer Stiefelhagen, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. 2016. Movieqa: Understanding stories in movies through question-answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4631–4640.
- Jesse Thomason, Subhashini Venugopalan, Sergio Guadarrama, Kate Saenko, and Raymond Mooney. 2014. Integrating language and vision to generate natural language descriptions of videos in the wild. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 1218–1227.
- Ramakrishna Vedantam, Xiao Lin, Tanmay Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. Learning common sense through visual abstraction. In *International Conference on Computer Vision*.
- Subhashini Venugopalan, Marcus Rohrbach, Jeffrey Donahue, Raymond Mooney, Trevor Darrell, and Kate Saenko. 2015. Sequence to sequence-video to text. In *Proceedings of the IEEE international conference on computer vision*, pages 4534–4542.
- Carl Vondrick, Deniz Oktay, Hamed Pirsiavash, and Antonio Torralba. 2016. Predicting motivations of actions by leveraging text. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2997–3005.
- Jason Weston, Antoine Bordes, Sumit Chopra, Alexander M Rush, Bart van Merriënboer, Armand Joulin, and Tomas Mikolov. 2015. Towards ai-complete question answering: A set of prerequisite toy tasks. *arXiv preprint arXiv:1502.05698*.
- Dejing Xu, Zhou Zhao, Jun Xiao, Fei Wu, Hanwang Zhang, Xiangnan He, and Yueting Zhuang. 2017. Video question answering via gradually refined attention over appearance and motion. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 1645–1653.

- Jun Xu, Tao Mei, Ting Yao, and Yong Rui. 2016. Msrvtt: A large video description dataset for bridging video and language. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 5288–5296.
- Hui Yang, Lekha Chaisorn, Yunlong Zhao, Shi-Yong Neo, and Tat-Seng Chua. 2003. Videoqa: question answering on news video. In *Proceedings of the eleventh ACM international conference on Multimedia*, pages 632–641. ACM.
- Yezhou Yang, Ching Teo, Hal Daumé III, and Yiannis Aloimonos. 2011. Corpus-guided sentence generation of natural images. In Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing, pages 444–454.
- Mark Yatskar, Luke Zettlemoyer, and Ali Farhadi. 2016. Situation recognition: Visual semantic role labeling for image understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5534–5542.
- Licheng Yu, Eunbyung Park, Alexander C Berg, and Tamara L Berg. 2015. Visual madlibs: Fill in the blank description generation and question answering. In *Proceedings of the ieee international conference on computer vision*, pages 2461–2469.
- Amir Zadeh, Michael Chan, Paul Pu Liang, Edmund Tong, and Louis-Philippe Morency. 2019. Social-iq: A question answering benchmark for artificial social intelligence. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8807–8817.
- Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. From recognition to cognition: Visual commonsense reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6720–6731.
- Luowei Zhou, Yingbo Zhou, Jason J Corso, Richard Socher, and Caiming Xiong. 2018. End-to-end dense video captioning with masked transformer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8739–8748.
- Linchao Zhu, Zhongwen Xu, Yi Yang, and Alexander G Hauptmann. 2017. Uncovering the temporal context for video question answering. *International Journal of Computer Vision*, 124(3):409–421.
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *Proceedings of the IEEE inter*national conference on computer vision, pages 19– 27.

Appendix

In this appendix, we provide the following supplementary information:

- Additional details about our dataset creation process, including statistics and analysis for V2C and V2C-QA datasets,
- Examples of commonsense descriptions generated by our V2C-Transformer model,
- Details about our human evaluation interface, protocol, and metrics.

A V2C Dataset Construction

Our dataset creation methodology is a three-step procedure as shown in Figure 9. In the first step, we use the caption to query ATOMIC (Sap et al., 2018) and retrieve the top-3 intentions, effects, and attributes. These are re-ranked by a BERT based model in the second step. The final step involves humans in the annotation process. We ask human annotators to select the most relevant descriptions, and to provide additional descriptions in their own words. The annotators also convert a subset of our dataset into complete sentence descriptions.

A.1 Querying from ATOMIC

For every video-caption pair in the MSR-VTT dataset, we select 3 most similar events from ATOMIC. These are then used to retrieve textual descriptions of three types – *intentions*, *effects*, *attributes* from ATOMIC.

A.2 BERT Ranking Model

We implement a Bidirectional Encoder Representations from Transformers (BERT) model (Devlin et al., 2019) as a ranking model to rank and retrieve *top-3* most plausible commonsense aspects to complement the ground truth caption. This is done by treating the ranking task as a binarized next sentence prediction (NSP) task, trained on the ATOMIC (Sap et al., 2018) dataset. When choosing the sentences A and B for each training pair, for 50% of the training pairs we choose the actual next sentence that follows A, and a random sentence from the ATOMIC as a negative sentence. This setting is consistent with the NSP task in (Devlin et al., 2019). We train our model in ATOMIC, and use it to expand video captions from MSR-VTT (Xu et al., 2016). Our BERT model consists of 12 transformer blocks, 12 attention heads, and 768 hidden dimensions (110M parameters in total). In total,



Figure 6: Next sentence prediction task in Bert model. A and B sentences are separated by token [SEP].

Commonsense Type	Accuracy (%)
Intention	84.87
Effect	86.53
Attribute	87.23
Average	86.21

Table 5: Accuracy of our BERT model for next sentence prediction on the ATOMIC test dataset split

we have 115,312 pairs for training/testing. We evaluate our model using accuracy of the prediction in the test set of ATOMIC which is 30% of the entire set. BERT can achieve 86.21% accuracy in NSP task on average as shown in Table 5. In addition, we also conduct human evaluations to measure the overall quality of the expanded V2C dataset (see "gold annotations" in Table. 3, main paper).

A.3 Human Labeling

With querying from ATOMIC and BERT reranking, we obtain commonsense descriptions that are relevant to the caption. However, we want to make sure that these descriptions are also relevant to the video. Thus we utilize human workers from Amazon Mechanical Turk (AMT) for selecting the most relevant commonsense descriptions. Our annotation interface is shown in Figure 10. We ask the annotators to select descriptions that are most relevant to the video and to the caption, and also encourage them to add their own commonsense descriptions. This makes our dataset more natural and human-like. This also allows us to remove noisy annotations that may be produced due to text-only ATOMIC querying. We show additional samples from our V2C dataset in Figure. 11, word cloud in Figure. 7 and word frequency in 8.

A.4 Benefits of the Three-Step Pipeline

Since our videos are annotated with captions, we use the captions to retrieve commonsense descriptions from ATOMIC. The ATOMIC dataset has comprehensive annotations for human activities, actions, and events and as such **covers most of the events in Msr-vtt**. Thus using these two datasets together is a natural step for creating our V2C dataset.



Figure 7: Word cloud figure of the intention commonsense annotations from our V2C dataset.

This purely caption-based retrieval unfortunately does not incorporate the latent aspects of the video, but only those from the caption. Moreover, since the video is not used for retrieving these, the commonsense annotations may be out-of-context. Thus, we bring in human annotators to watch the video, read the caption, and then use the set of descriptions from ATOMIC to select the relevant once and to discard the irrelevant or out of context descriptions. The human annotators then provide annotations about intention, effect, and attribute in their own words. The ATOMIC retrieved descriptions help the human annotators to get an idea about the task and also get a glimpse of the format of the desired annotations. This significantly reduces the noise in human annotations.

To guarantee and measure the overall quality of our V2C dataset, we have conducted human evaluations on the V2C annotations. Our results shows that 86.29% of the \langle video-caption-commonsense \rangle triples are labeled as reasonable samples (see "Gold Annotations" in main paper, Table. 3), **verifying the quality** of our dataset

B V2C-QA Dataset

For the V2C Question Answering task, we repurpose our V2C dataset and convert it to a question-answering dataset. We choose a subset of 1500 videos: 1250 for training and 250 for testing, following the same train-test split as MSR-VTT. We use SpaCy linguistic features (Honnibal and Montani, 2017) along with the LemmInflect library (https://github.com/bjascob/LemmInflect) and template-based generation to convert the captions, intentions, effects, and attributes from V2C to create questions and ground-truth answers. Our templates are linguistically diverse, natural, and grammatically sound. We have 21 types of templates with each template having numerous possibilities for combi-

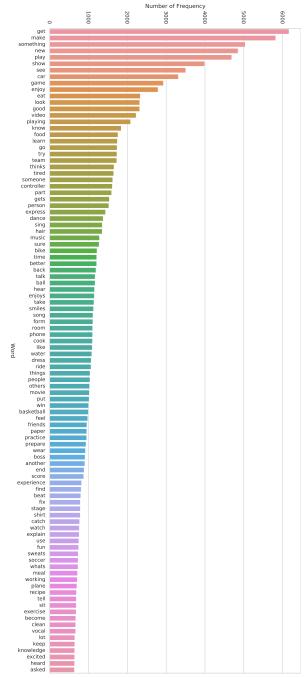


Figure 8: Top-100 most frequent words in our V2C dataset (stop words are ignored).

nations of the slots in the template. Thus we get 21 types of questions (7 each for intention, effect, and attribute) as shown in Table 6. Since our task is open-ended question-answering, our questions are annotated with *all possible* correct answers for that question. To get answers for the "negative" questions as shown in Table 6, we use the adversarial matching strategy similar to (Zellers et al., 2019), by using RoBERTa (Liu et al., 2019) similarity. We will release our V2C-QA question and

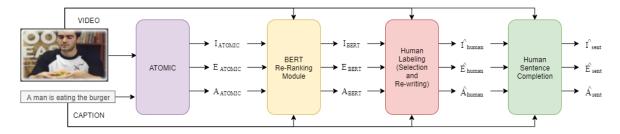


Figure 9: The data creation flow for V2C. We use the retrieved videos and captions from MSR-VTT and use the BERT re-ranking module to obtain a list of top-3 intentions (I), effects (E), and attributes (A). These are then further improved by human labeling. A subset of annotations is also converted to full sentences by human annotators.

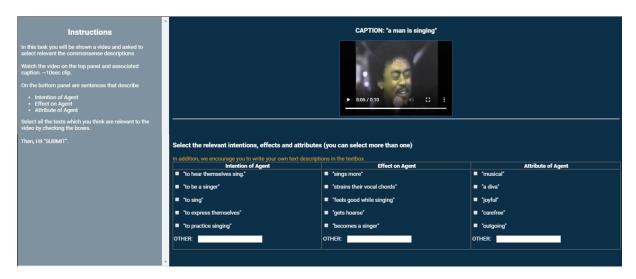


Figure 10: Our human labeling interface. We ask human workers to select relevant commonsense descriptions as well provide additional texts in their own words

answer generation code publicly.

C Qualitative Generation Results

We show additional **V2C-Completion** samples by our V2C-Transformer model in Table. 7.

D Human Evaluation

Human evaluation is one of the important part to verify the performances of our model and the quality of the V2C dataset. In this section we describe our setup for human evaluation of the captions and commonsense descriptions in our dataset as well as those generated by our models.

D.1 Amazon Mechanical Turk Interface

We conduct our human evaluations by crowdsourcing ratings from workers on Amazon Mechanical Turk (AMT). We do these human evaluations on the same test set used for our automated metrics. We show an example of our interface in Figure 12 and 13 which shows the screenshot of the rating task as

seen by the workers. The workers are given explicit instructions about this rating task, and depending on the task are asked to rate the commonsense descriptions and the caption.

For the V2C-Completion task, the workers are provided with the video and the ground-truth caption and asked to rate the only the generated commonsense (intention, effect or attribute) on a scale of 1 to 5. The workers are asked to provide this rating on the basis of whether the generated text is relevant to the video, i.e whether the caption/commonsense can plausibly complete the given event.

For the V2C-Generation task, the workers are asked to rate the caption as well as the commonsense texts with respect to the video. The workers are also asked to conduct identical tasks for the gold (ground-truth annotations) in our new V2C dataset.



Question Type	Question	Answer
Intention Intention (Negative) Intention (Action) Intention (Action, Negative) Intention (Why) Intention (Yes-No) Intention (Yes-No, Negative)	What might be the goal of the person? What could the person not want to achieve? What prompts the person to do the action? What did not lead the person to act like that? Why might the person be doing the action? Does the person wish to express himself? Does the person want to not get recognition?	to record a music video to bake a cake to express themselves to feed the dog to entertain viewers Yes No
Effect Effect (Negative) Effect (Action) Effect (Action, Negative) Effect (How) Effect (Yes-No) Effect (Yes-No, Negative)	What will the person do after this? What does not happen as a result? What does the dancing end up in? What will not happen due to the action? How does the person feel after performing? Could the person put it on YouTube as a result? Will the person not learn a new dance?	puts the video on YouTube the person gets sad becomes tired feels tense feels accomplised Yes No
Attribute Attribute (Negative) Attribute (How) Attribute (Action, How) Attribute (Yes-No, Action) Attribute (Yes-No) Attribute (Yes-No, Negative)	What trait does the man possess? What attribute does not match with the person? How can the person be described? How can the dancing person be characterized? Is the person who is singing smiling? Is the person entertaining? Is the person not tense?	musical angry entertaining rhythmic Yes Yes Yes

Table 6: Examples of open-ended V2C-QA samples

D.2 Scheme for Validity

Our ratings are measured on a scale of 1 to 5. Annotations which receive a score greater than 3 are considered "valid", so as to be consistent with the binary ratings used by (Bosselut et al., 2019) for their experiments. We then compute average validity scores for each commonsense aspect: intention, attribute and effect.

D.3 Statistics of Human Evaluations

In order to further analyze the human evaluations on our generated outputs, we use three metrics - standard deviation of the ratings, inter-rater agreement score (IRAS) and a smooth version of IRAS. Standard Deviation was calculated per sample based on the evaluations provided by multiple workers on each sample. We do so to evaluate how consistent our AMT workers are and how much they deviate or agree with each other. We use three different metrics so as to analyze our data and generations through multiple lenses, to be certain that the outputs and annotations are high-quality.

D.3.1 Inter-Rater Agreement Score

Inter-Rater Agreement Score is computed as the average of the percentage of raters for each sample that agree with the majority opinion. Let m be the size of the test-set, and n be the number of rating. Let $\mathcal{R}^j = \{r_1, \ldots, r_n\}$ be the set of ratings for test sample j. Then the mode r_{mode} is defined as the most frequently occurring (majority) rating in the set of ratings \mathcal{R}^j , i.e. $r_{mode} = \text{MODE}(\mathcal{R}^j)$.

Inter-Rater Agreement Score IRAS $_{agree}$ is the average percentage of raters that agree with the majority opinion r_{mode} :

$$\mathbf{IRAS} = 100 \times \sum_{i=1}^{m} \frac{\sum_{i=1}^{n} \mathbb{I}(r_i = r_{mode})}{n \times m}, \quad (5)$$

where \mathbb{I} is the indicator function.

D.3.2 Smooth Inter-Rater Agreement Score

While IRAS acts as a good metric to find out how our dataset fares in terms of rater agreement, it suffers from a flaw. Irrespective of the value of ratings, the indicator function \mathbb{I} returns 0 for the tuple of ratings (1,5) as well as (4,5), although the

Intention	Caption	Effect	Attribute
to entertain people	a band is performing for a crowd	gets applause	acting
to try out PersonY's new	a man checks out detail on a car	gets a speeding ticket	helpful
car			
to learn about current	a complex news host gives an update	gets informed about cur-	talkative
events	on rappers.	rent political events	
to be in a good mood	a group of people trying to perform an exorcism on a girl	gets applause	fun
to show his knowledge- able	there is an old man is answering to somebody questions	gets another question	sporty
to score a point	a man is shooting a basketball ground	gets exercise	helpful
to share their message	a man giving a speech to important people	gets applause	orator
to be safe from anything	a group of people are being chased by	gets tired from taking	scared
that lurks in the dark	crocodiles	pictures	
to be informed about the world	a girl is describing about hot news	learns about whats hap- pening worldwide	gossipy
to watch something interesting	a children s television show clip	smiles at the screen	entertained
to enjoy the evening with	a band composed of older gentlemen	gets tired form dancing	fun
the concert band	are playing blue grass music on a small stage and people are dancing along to the music swing-style		
to be part of the team	there is a woman playing badminton in a court	gets tired after exercise	athletic
to try out person ys new	a boy explaining the features of a car	they check car websites	helpful
car		online to look at deals	•
to escape reality	a man explaining a video game	takes the video game home	gamer
to cook something	there is a man in black cutting the green leaves on the desk	gets clean dishes	hungry

Table 7: Illustrative samples generated by our V2C-Transformer model on V2C-completion task.

	Type	Std. Dev $(\%) \downarrow$		IRA	AS (%) ↑	smooth-IRAS (%) ↑		
		ATTENCDEC	V2C-Transformer	ATTENCDEC	V2C-Transformer	ATTENCDEC	V2C-Transformer	
	Intention	17.99	15.02	56.02	59.80	69.43	73.36	
V2C-Completion	Effect	19.63	18.39	58.03	56.76	69.28	69.47	
	Attribute	10.54	9.74	69.06	71.28	80.24	81.83	
	Average	16.05	14.38	61.04	62.61	72.98	74.89	
	Intention	17.60	16.27	57.84	58.47	70.66	72.10	
V2C-Generation	Effect	18.54	17.56	56.69	57.40	69.54	70.21	
	Attribute	15.42	13.16	59.80	62.25	73.51	76.12	
	Average	17.19	15.66	58.11	59.37	71.24	72.81	

Table 8: A comparison of the statistics of human evaluation scores for both tasks using the baseline (ATTENCDEC model vs. our model (V2C-Transformer)

ratings of 4 and 5 are close to each other but 1 and 5 are opposite. So to avoid this, we replace the indicator function with a smooth exponential term. The smooth inter-rater agreement score is given by:

$$\mathbf{IRAS}_{smooth} = 100 \times \sum_{j=1}^{m} \frac{\sum_{i=1}^{n} \left(\frac{1}{2}\right)^{|r_i - r_{mode}|}}{n \times m}.$$
(6)

D.3.3 Results

Table 8 shows our analysis in terms of the three metrics described above. Our V2C-Transformer architecture consistently outperforms the baseline model ATTENCDEC (Gao et al., 2017) in all three metrics for each type of commonsense. This means that raters are more consistent with their ratings (in terms of deviation or agreement) for commonsense descriptions generated by our model.

The person wants:

As an effect, the person: The person is:

to express themselves to sing a song to make life more pleasant



put it on YouTube learns a new dance gets into their rhythm outgoing enthusiastic energetic

Caption

A guy sings a song in a music video

Rewritten Story:

Because he wants to express himself, a guy sings a song in a music video, and he will upload it to YouTube soon. He is quite an enthusiastic guy.

to show off at the gym to acquire new footwear to wear appealing clothes



becomes obsessed gets their picture taken grabs attention from people stylish trendy fashionable

Caption

Girls trying on new sports bra.

Rewritten Story:

In order to purchase new sportswear, girls trying on new sports bra, and they may grab attention from people later. They are all stylish person.

to win the race to earn a medal to win the competition



runs a race is congratulated at the finish line focuses on the race athletic competitive

Caption

Groups of runner get prepared to run a race.

Rewritten Story:

Since the athletes are trying to win the race, groups of runner get prepared to run a race, and they will run and get congratulated at the finish line soon. They are athletic.

to show appreciation to be accommodating to talk to them



sweats from nervousness shares information communicates empathetic talkative conversational

Caption President Obama calls a team to congratulate them.

Rewritten Story:

To show his appreciation to the winners, President Obama calls a team to congratulate them, the girls will got sweats because of that. The Obama is so talkative.

to have a conversation convey information to give speech



Caption A group of males speaking to each other at a meeting.

gives a rebuttal gets to meet the host loses their voice due to loud talking/yelling extroverted polite speaker

Rewritten Story:

In order to convey with each other the information, a groups of males speaking to each other at a meeting, they will get into a rebuttal soon. The people have the attribute to be extroverted.

to get to her destination to get somewhere to drive fast



travels to a different city arrives at their destination enjoy driving traveling a good driver speedy

Caption

A man drives a vehicle through the countryside.

Rewritten Story:

To get to his destination as soon as possible, a man drives a vehicle through the countryside, he may soon arrives at his destination. The man is a good driver.

to get the computer working set up system to clean the viruses from his computer



turns off the computer boots up the computer spends money on a new computer

busy smart informative

Caption A woman in a business suit looking at a computer monitor.

Rewritten Story:

Because the computer is not working and the woman is trying to fix it, a woman in a business suit looking at a computer monitor, she will boots the computer first soon. She is a very informative person.

Figure 11: Qualitative examples of our V2C dataset.

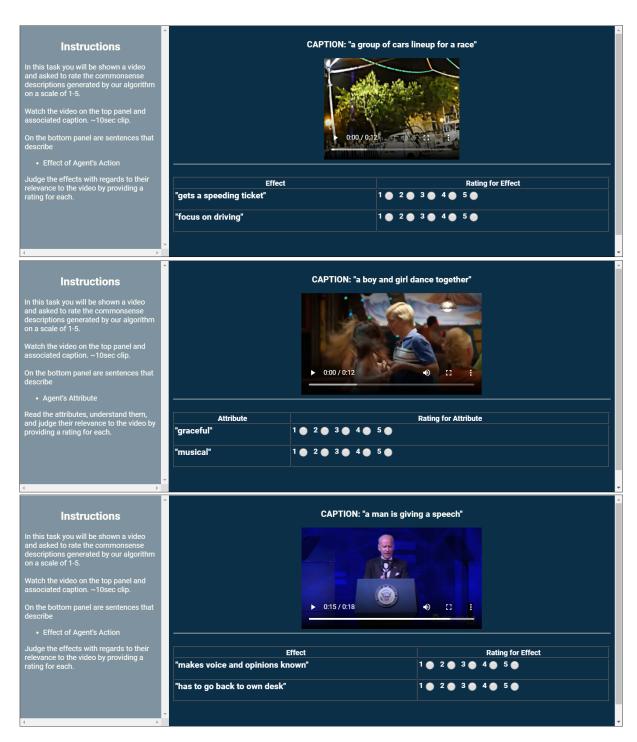


Figure 12: Snapshot of our AMT human evaluation interface for V2C-completion task.



Figure 13: Snapshot of our AMT human evaluation interface for V2C-generation task.

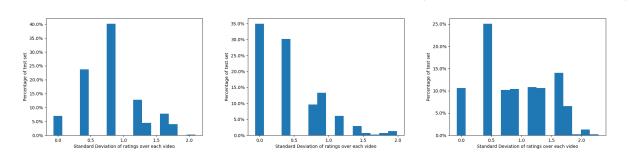


Figure 14: V2C-Completion task using the ATTENCDEC model.

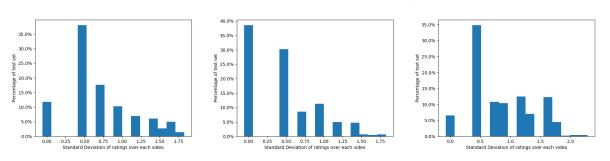


Figure 15: V2C-Completion task using our V2C-Transformer model.

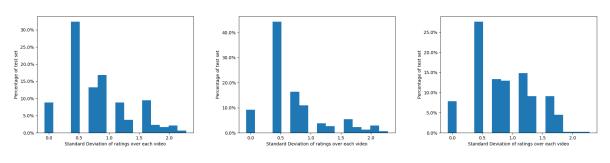


Figure 16: V2C-Completion task using the ATTENCDEC model.

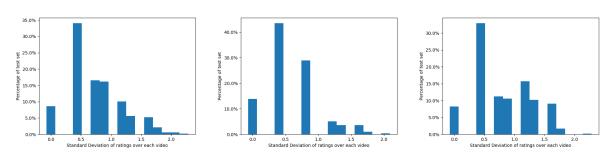


Figure 17: V2C-Generation task using our V2C-Transformer model.

Figure 18: Standard deviation histograms of human ratings across models and split (*From left to right: Intention, Attribute, Effect*). X-axis denotes standard deviation value and Y-axis denotes percentage of test set samples.