

Vision Guided Generative Pre-trained Language Models for Multimodal Abstractive Summarization

Tiezheng Yu*, Wenliang Dai*, Zihan Liu, Pascale Fung

Center for Artificial Intelligence Research (CAiRE)

Department of Electronic and Computer Engineering

The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong

{tyuah,wdaiai,zliucr}@connect.ust.hk, pascale@ece.ust.hk

Abstract

Multimodal abstractive summarization (MAS) models that summarize videos (vision modality) and their corresponding transcripts (text modality) are able to extract the essential information from massive multimodal data on the Internet. Recently, large-scale generative pre-trained language models (GPLMs) have been shown to be effective in text generation tasks. However, existing MAS models cannot leverage GPLMs' powerful generation ability. To fill this research gap, we aim to study two research questions: 1) how to inject visual information into GPLMs without hurting their generation ability; and 2) where is the optimal place in GPLMs to inject the visual information? In this paper, we present a simple yet effective method to construct vision guided (VG) GPLMs for the MAS task using attention-based add-on layers to incorporate visual information while maintaining their original text generation ability. Results show that our best model significantly surpasses the prior state-of-the-art model by 5.7 ROUGE-1, 5.3 ROUGE-2, and 5.1 ROUGE-L scores on the How2 dataset (Sanabria et al., 2018), and our visual guidance method contributes 83.6% of the overall improvement. Furthermore, we conduct thorough ablation studies to analyze the effectiveness of various modality fusion methods and fusion locations. ¹

1 Introduction

Multimodal abstractive summarization (MAS) aims to take advantage of data from multiple modalities and provides a short, concise and readable textual summary to let users quickly acquire their essential information (Sanabria et al., 2018; Palaskar et al., 2019; Liu et al., 2020). MAS has become an increasingly popular research area thanks to the proliferation of online multimedia content and the increasing availability of multimodal data.

* The two authors contribute equally.

¹The code is available at: <https://github.com/HLTCHKUST/VG-GPLMs>

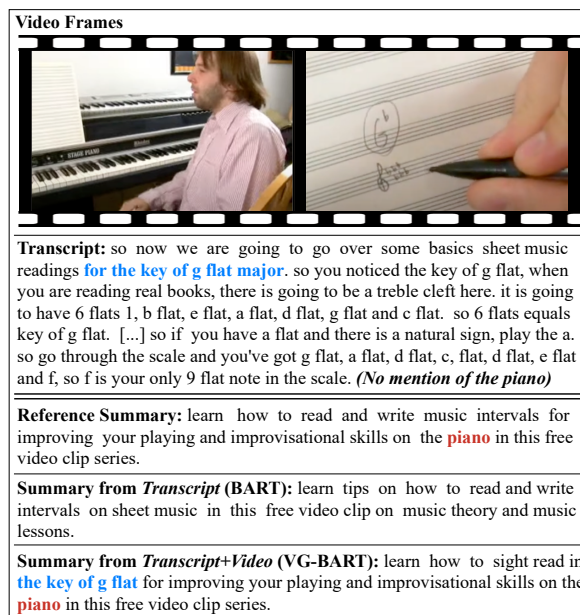


Figure 1: An example of MAS. As input data, we show two representative video frames and the transcript, with [...] representing omitted unimportant text. As illustrated, some information is emphasized (e.g. **the key of g flat**) or only exists (e.g. **piano**) in the visual signal. We also compare the human-generated reference summary and our model-generated summaries with/without video frames in the input data.

As illustrated in Figure 1, the MAS models need to generate a concise summary by effectively utilizing two modalities: a video and its transcript. Therefore, we emphasize that leveraging a powerful text generation model and an effective combination of the vision and text modalities are key to constructing good MAS models. Recently, Transformer-based (Vaswani et al., 2017b) sequence-to-sequence (Seq2Seq) large-scale generative pre-trained language models (GPLMs), such as BART (Lewis et al., 2019), T5 (Raffel et al., 2019), PEGASUS (Zhang et al., 2020a) and ProphetNet (Qi et al., 2020), have shown remarkable performance on text generation tasks, including abstractive text summarization. However, leveraging and adapting GPLMs to MAS

is still an unexplored research direction. To explore this direction, two main questions need to be answered: Firstly, *how* can we inject visual information into the text-only GPLMs so that the models can understand both modalities and allow cross-modal interactions, and more importantly, how can this injection operation be conducted without damaging GPLMs’ original text generation ability? Secondly, *where* is the optimal place in GPLMs to inject the visual information? This needs to be explored, as there are many sub-layers in the encoder and decoder of GPLMs and a sub-optimal location might result in unsatisfactory performance.

In this paper, to fill the research gap, we present a simple yet very effective method to construct vision guided (VG) GPLMs (VG-BART and VG-T5) for the MAS task. Specifically, to answer the first of the aforementioned questions, we insert attention-based add-on layers to GPLMs to incorporate visual information without modifying the original architecture. In this way, all the pre-trained model weights can be used during fine-tuning so as to preserve their original text generation ability. We try with two types of attention mechanisms for the text-vision fusion and interaction: 1) Cross-modal Dot-product Attention; and 2) Cross-modal Multi-head Attention. Moreover, we also investigate the effects of using a forget gate and a visual transformer encoder along with the attention mechanisms. To answer the second question, we enumerate almost all possible locations in GPLMs for injecting add-on layers, and show a thorough comparison and analysis in Section 5. We evaluate our models on the How2 dataset (Sanabria et al., 2018). Experimental results demonstrate that our best model surpasses the prior state-of-the-art model by 5.7 ROUGE-1, 5.3 ROUGE-2, and 5.1 ROUGE-L scores. To ensure this improvement does not purely come from the GPLMs, we also evaluate the corresponding text-only model, and the results show that the injected visual guidance contributes 83.6% of the overall improvement on average of all ROUGE scores.

Our contributions in this work are threefold:

- To the best of our knowledge, we are the first to inject visual information into text-only GPLMs, and to use it for the MAS task.
- We systematically study two research questions: 1) how to inject visual information into GPLMs without hurting their generation ability; and 2) where is the optimal place in GPLMs to inject the visual information?

- Our model significantly outperforms the state-of-the-art model on the How2 dataset, and the injected visual guidance contributes 83.6% of the overall improvement.

2 Related Work

2.1 Abstractive Text Summarization

Abstractive text summarization aims to generate short, concise and readable text that can capture the most salient information of the input documents. Thanks to the Seq2Seq framework (Sutskever et al., 2014) and attention mechanisms, deep neural networks have achieved remarkable results on summarization tasks (Paulus et al., 2017; Zhang et al., 2020b; Yu et al., 2021). Recently, GPLMs (Lewis et al., 2019; Raffel et al., 2019; Zhang et al., 2020a; Qi et al., 2020) have been widely used in abstractive text summarization and have achieved start-of-the-art performance. The most significant difference between abstractive text summarization and multimodal abstractive summarization lies in whether the input contains data of more than one modality.

2.2 Multimodal Abstractive Summarization

Recently, many studies have been performed on multimodal learning (Mroueh et al., 2015; Antol et al., 2015; Donahue et al., 2015; Zadeh et al., 2017; Dai et al., 2020, 2021). However, only a few have investigated MAS. Li et al. (2017) collected a multimodal corpus of news articles containing 500 videos of English news articles paired with human-annotated summaries. Sanabria et al. (2018) introduced the How2 dataset, which contains about 2,000 hours of short instructional videos, each coming with a summary of two to three sentences. Palaskar et al. (2019) proposed a multi-source Seq2Seq model with hierarchical attention to integrate information from different modalities into a coherent summary. Meanwhile, Liu et al. (2020) proposed a multi-stage fusion network with the fusion forget gate module, which can model the fine-grained interactions between multi-source modalities. To the best of our knowledge, no previous work has leveraged GPLMs’ generation ability to tackle the MAS task, and we are the first to systematically study multiple multimodal fusion methods based on GPLMs.

2.3 Vision-Language Large Pre-trained Transformer Models

With the remarkable success of large-scale unsupervised pre-training in NLP (Devlin et al., 2019;

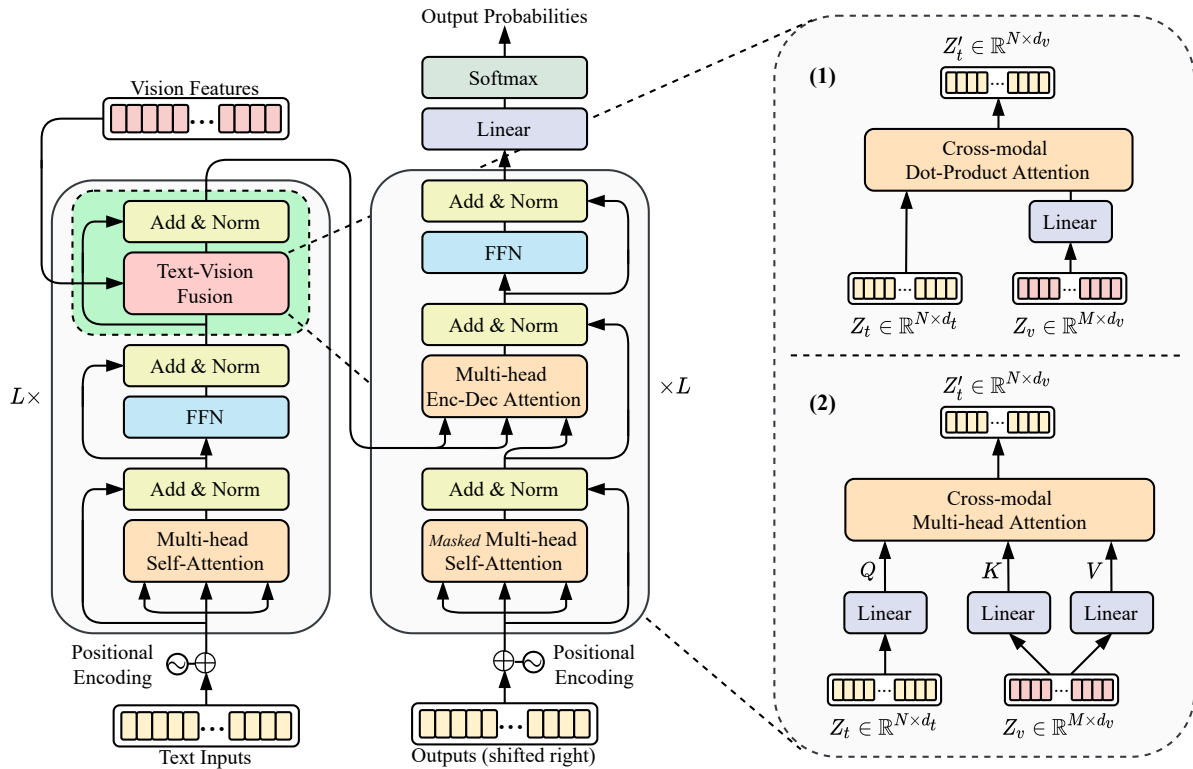


Figure 2: An overview of our proposed VG GPLMs. It is built based on the Transformer-based Seq2Seq GPLMs (left). To inject visual information, we insert add-on sub-layers (the green dashed block) by mainly leveraging two kinds of attention-based text-vision fusion mechanism (right): 1) Cross-modal Dot-Product Attention; and 2) Cross-modal Multi-head Attention. Although we draw the add-on sub-layers in the encoder, they can also be placed in the decoder in a similar way. We compare the effects of different injection locations in Section 5.

Liu et al., 2019; Radford et al., 2019), pre-training large vision-language (VL) models has also become more and more popular in recent years. Rather than designing task-specific architectures, pre-training results in a general backbone model by feeding it with a large amount of data and then fine-tune it to different downstream tasks. Among the current VL pre-training work, most has been focusing on VL understanding by training BERT-style Transformer models (Sun et al., 2019; Tan and Bansal, 2019; Su et al., 2020; Li et al., 2020; Chen et al., 2020) and finetune them on various VL classification tasks (Goyal et al., 2017; Zellers et al., 2019; Suhr et al., 2019). These models usually receive a pair of text and image as input, where the image is processed into objects (Zhang et al., 2021), patches (Kim et al., 2021), or pixels (Huang et al., 2020) before feeding into the VL model. For VL text generation, Zhou et al. (2020) presented a model for both visual question answering and image captioning (Chen et al., 2015). Additionally, Cho et al. (2021) introduced an encoder-decoder Transformer model that unifies all VL tasks as

generative tasks. Although prior work has made much progress on VL pre-training, the problem of generating text given text and video input (E.g. the How2 dataset) is not well studied under the VL pre-training setting, except by Luo et al. (2020), who proposed a dual-stream model for both VL classification and generation with video data. However, compared to GPLMs in NLP such as BART (Lewis et al., 2019) and T5 (Raffel et al., 2019), their text generation ability is limited as the training data is much smaller.

In this paper, we propose to tackle VL tasks and utilize the advantage of pre-training from a different angle by inserting add-on layers to the text-only GPLMs and fine-tuning them on multimodal tasks to incorporate visual information. This takes advantage of GPLMs’ superior generation ability to generate vision-aware texts. Of the very few works that have also considered this direction, Rahman et al. (2020) proposed the multimodal adaptation gate, which fuses data of other modalities to the textual embeddings in BERT. However, their method requires all modalities to have the same sequence

length, which is rare for most datasets. Additionally, they only attempted to address the sentiment analysis task and did not explore text generation.

3 Vision Guided GPLMs

To take advantage of the superior text generation ability of the text-only Seq2seq GPLMs and adapt them to the MAS task, we present Vision guided (VG) GPLMs. Specifically, we leverage BART (Lewis et al., 2019) and T5 (Raffel et al., 2019) to construct VG-BART and VG-T5.

In this section, we start by revisiting the text-only Seq2seq GPLMs in Section 3.1. These serve as the backbone of our proposed model and also one of the baselines. Then, we discuss the approach for extracting visual features from video clips in Section 3.2, as well as how to further process them. Finally, in Section 3.3, we introduce two types of text-vision fusion mechanism to guide the GPLMs to generate vision-aware summaries.

3.1 Overview of GPLMs for Summarization

Transformer-based (Vaswani et al., 2017b) Seq2Seq GPLMs generalize architectures like BERT (Devlin et al., 2019) and GPT (Radford et al., 2018) by including a bi-directional encoder and a uni-directional (left-to-right) decoder. An overview of this architecture is depicted on the left side of Figure 2 (except the green dashed block).

At the entry of the GPLM, the input text is first tokenized and converted to a sequence of token embeddings $X_t \in \mathbb{R}^{N \times d_t}$, in which N is the sequence length and d_t is the feature dimension. To retain the positional information, positional encodings (Vaswani et al., 2017a) $E_{pe} \in \mathbb{R}^{N \times d_t}$ are added to the token embeddings pointwisely (Eq. 1), which forms the input features Z_0^{enc} to the encoder.

$$Z_0^{enc} = X_t + E_{pe} \quad (1)$$

As illustrated in Figure 2, the encoder is composed of a stack of L encoder layers, each containing two sub-layers: 1) Multi-head Self-Attention (MSA, Eq. 2) and 2) Feed-Forward Network (FFN, Eq. 3). In addition, after each sub-layer, there is a residual connection (He et al., 2015; Wang et al., 2019) followed by a layer normalization (LN) (Ba et al., 2016). See Appendix A and B for more details of the MSA and FFN.

$$Z_l^{enc'} = \text{LN}(\text{MSA}(Z_{l-1}^{enc}) + Z_{l-1}^{enc}) \quad (2)$$

$$Z_l^{enc} = \text{LN}(\text{FFN}(Z_l^{enc'}) + Z_l^{enc'}) \quad (3)$$

Similar to the encoder, the decoder also consists of a stack of L decoder layers, but with two differences. Firstly, the MSA is masked to prevent positions from attending to subsequent positions (keep the decoder in a left-to-right direction). Secondly, there is one more multi-head encoder-decoder attention sub-layer, which uses the decoder embeddings to attend over the output embeddings of the encoder to incorporate the encoded information.

Specifically, in our experiments, we adopt the pre-trained BART (Lewis et al., 2019) and T5 (Raffel et al., 2019), which both follow this architecture with different training schemes. To fine-tune them on the abstractive text summarization task, the input to the encoder is the article or transcript, and the decoder learns to generate the summaries.

3.2 Video Feature Extraction

For each video clip, following previous works (Sanabria et al., 2018; Palaskar et al., 2019; Khullar and Arora, 2020), a 2048-dimensional feature representation is extracted for every 16 non-overlapping frames using a 3D ResNeXt-101 model (Hara et al., 2018), which is pre-trained on the Kinetics dataset (Kay et al., 2017). Therefore, each data sample will have a sequence of 2048- d vision feature vectors of length M . These features can be used directly as the visual input to the text-vision fusion mechanism.

In addition, in order to better model the intramodal dynamics and enhance the vision specific temporal information, we further process the extracted sequence of visual features using a Transformer (Vaswani et al., 2017a) encoder (VTF) with positional encodings. Experiments illustrate that this additional encoding process can further boost the performance of our model (Section 5).

3.3 Text-vision Fusion

As exhibited in Figure 2, we insert a third sub-layer (the green dashed block) into each encoder layer, which contains the text-vision fusion mechanism and also a residual connection followed by a layer normalization. We propose two types of text-vision fusion mechanism, as shown on the right-hand side of the figure. Given the textual input $Z_t \in \mathbb{R}^{N \times d_t}$ and visual input $Z_v \in \mathbb{R}^{M \times d_t}$, the fusion mechanism produces vision guided output $Z'_t \in \mathbb{R}^{N \times d_t}$ that has a same dimension as the textual input, which allows the continual stacking of layers.

Dot-product Attention Based Fusion. Before performing dot-product attention between the textual and visual features, we first project the visual features Z_v to the same dimensional space as the textual features (Eq. 4). Then, we calculate the dot-product and apply the softmax function to get the attention score matrix A (Eq. 5). Finally, the input textual features Z_t are concatenated with the attention weighted visual features AZ_v and then projected by another linear transformation to output the vision guided textual features Z'_t (Eq. 6).

$$Z'_v = Z_v W_1, Z'_v \in \mathbb{R}^{M \times d_t} \quad (4)$$

$$A = \text{Softmax}(Z_t Z'_v{}^T), A \in \mathbb{R}^{N \times M} \quad (5)$$

$$Z'_t = \text{Concat}(Z_t, AZ_v) W_2 \quad (6)$$

Additionally, we build a variant of this fusion, which uses the linearly transformed visual features AZ'_v for the concatenation in Eq. 6 instead of the original AZ_v . A comparison of their performance is shown in Section 5.

Multi-head Attention Based Fusion. Inspired by prior works (Yu et al., 2019; Tsai et al., 2019), we propose a vision guided multi-head attention mechanism for the text-vision fusion. The query Q is linearly projected from the input textual features, and the key K and value V are linearly projected from the visual features (Eq. 7 - 9). Then, a cross-modal multi-head attention (CMA) is applied to get the text queried visual features O (Eq. 10). Finally, we obtain the vision guided output Z'_t by concatenating the input textual features Z_t and O , and linearly project it to the desired dimension (Eq. 11).

$$Q = Z_t W_q, Q \in \mathbb{R}^{N \times d_c} \quad (7)$$

$$K = Z_v W_k, K \in \mathbb{R}^{M \times d_c} \quad (8)$$

$$V = Z_v W_v, V \in \mathbb{R}^{M \times d_c} \quad (9)$$

$$O = \text{CMA}(Q, K, V), O \in \mathbb{R}^{N \times d_c} \quad (10)$$

$$Z'_t = \text{Concat}(Z_t, O) W_3 \quad (11)$$

In addition, we also explore the effects of using a forget gate (Liu et al., 2020) in the text-vision fusion. Given the CMA output $O \in \mathbb{R}^{N \times d_c}$ in Eq. 10, we construct a forget gate mask $F \in \mathbb{R}^{N \times d_c}$ (Eq. 12) and do a point-wise multiplication with O to output the updated O' (Eq. 13).

$$F = \text{Sigmoid}(\text{Concat}(O, Z_t) W_f) \quad (12)$$

$$O' = F \otimes O \quad (13)$$

The forget gate can potentially remove redundant and noisy information from the video features, which also helps the model to learn to discard needless visual information to retain its pre-trained text generation ability.

4 Experimental Settings

4.1 How2 Dataset

How2 (Sanabria et al., 2018) is a large-scale dataset of open-domain videos, covering 22 different topics such as cooking, exercise, yoga and music. It consists of 79,114 short instructional videos (73,993 for training, 2,965 for validation and 2,156 for testing). Each video is accompanied by a human-generated transcript and a short text summary. At the word level, the average lengths of transcripts and summaries are 291 and 33, respectively.

4.2 Implementation Details

Data pre-processing. We pre-process the transcripts data by truncating or padding them into sequences of 512 tokens after tokenization. For the videos, after the feature extraction as described in Section 3.2, we also truncate or pad the sequence length to 256.

Hyper-parameters. We use BART-base and T5-base as the pre-trained GPLMs to construct VG-BART and VG-T5, in which $L = 6$ for both encoder and decoder. For the VTF mentioned in Section 3.2, we use a 4-layer encoder with 8 attention heads and a 2048 feed-forward dimension. In the decoding stage, we use beam search with a beam size of 5. The decoding process will not stop until an end-of-sequence (EOS) token is emitted or the length of the generated summary reaches to 64 tokens. Following Lewis et al. (2019) and Raffel et al. (2019), we use learning rates $6e^{-4}$ and $3e^{-5}$ to fine-tune the pre-trained parts of model weights. While for the newly added layers, we set the learning rate to $1.5e^{-4}$. For all of our experiments, we use a batch size of 120.

Optimizer. During training, we use the Adam optimizer (Kingma and Ba, 2015) with $\beta_1 = 0.9$, $\beta_2 = 0.999$ and a weight decay of $1e^{-5}$. Additionally, we apply a scheduler to decay the learning rate to 95% of the current one after every 10 epochs. We train all the models for 60 epochs with an early stop of 5 using the ROUGE-2 score (Xiao and Carenini, 2019) on the validation set.

Input	Method	R-1	R-2	R-L	B-1	B-2	B-3	B-4	M	C	CF
Transcript	S2S*	58.6	40.6	53.8	55.2	45.6	39.9	35.8	27.6	2.35	-
	PG*	57.2	39.5	52.8	55.3	45.6	39.8	35.7	26.8	2.13	-
	TF*	59.0	41.0	54.3	56.6	46.7	40.8	36.6	27.7	2.30	-
	T5	62.8	45.0	57.5	60.5	50.4	44.2	39.6	30.6	2.76	61.7
	BART	64.0	46.4	58.9	62.4	52.6	46.4	42.0	31.7	2.97	63.9
Transcript +Video	HA (RNN)*	60.3	42.5	55.7	57.2	47.7	41.8	37.5	28.8	2.48	-
	HA (TF)*	60.2	43.1	55.9	58.6	48.3	43.3	38.1	28.9	2.51	-
	MFFG (RNN) [†] *	62.3	46.1	58.2	59.1	50.4	45.1	41.1	30.1	2.69	-
	MFFG (TF)*	61.6	45.1	57.4	60.0	50.9	45.3	41.3	29.9	2.67	-
	VG-T5 (Dot-product)	63.0	44.9	57.6	60.1	49.8	43.4	38.8	30.3	2.74	61.4
	VG-T5 (Multi-head)	63.3	45.3	58.0	60.7	50.8	44.7	40.2	31.0	2.86	62.8
	VG-BART (Dot-product)	66.1	49.3	61.2	64.5	55.1	49.2	44.8	33.2	3.18	66.9
	VG-BART (Multi-head)	66.3	49.4	61.4	64.1	54.8	48.9	44.6	33.1	3.18	67.3

Table 1: Evaluation results of baselines and our proposed models on the How2 dataset. We compare the performance of using transcript only and transcript+video. The † indicates the previous state-of-the-art model. Results with * mark are taken from the previous work (Liu et al., 2020). We denote ROUGE, BLEU, METEOR, CIDEr and Content F1 by R, B, M, C and CF respectively.

Input	Method	R-1	R-2	R-L	B-1	B-2	B-3	B-4	M	C	CF
Transcript +Video	VG-BART (Multi-head)	66.3	49.4	61.4	64.1	54.8	48.9	44.6	33.1	3.18	67.3
	w/ FG	67.3	50.7	62.4	65.0	55.9	50.1	45.7	33.8	3.25	72.5
	w/ VTF	67.3	50.9	62.6	64.9	56.0	50.1	45.7	33.7	3.20	72.1
	w/ FG+VTF	68.0	51.4	63.3	65.2	56.3	50.4	46.0	34.0	3.28	69.7

Table 2: Further Evaluation of adding forget gate (FG) and visual transformer encoder (VTF) to our best model setting in Table 1 on the How2 dataset. VG-BART+FG+VTF largely surpasses the previous state-of-the-art model.

Software and hardware. We use the deep learning framework PyTorch (Paszke et al., 2019) to implement our code and PyTorch-Lightning² for the distributed training. We use four Nvidia GeForce RTX 2080 Ti GPUs for all of our experiment.

4.3 Baselines

Apart from the text-only GPLMs BART (Lewis et al., 2019) and T5 (Raffel et al., 2019), we use the following baselines to compare with our proposed models, including simple models that only accept text input, as well as prior state-of-the-art models that accept text and vision modalities.

S2S (Luong et al., 2015). S2S is a standard Seq2seq model that uses RNNs for both encoder and decoder with a global attention mechanism (Bahdanau et al., 2014).

PG (See et al., 2017). The pointer generator (PG) network augments S2S by having a copy module

²<https://github.com/PyTorchLightning/pytorch-lightning>

to reproduce key information accurately as well as mitigating the out-of-vocabulary issue.

TF (Vaswani et al., 2017b). TF is the standard Transformer-based Seq2seq model, which proposes the novel multi-head attention mechanism.

HA (RNN/Transformer) (Palaskar et al., 2019). A multi-source Seq2seq model with hierarchical attention (HA) (Libovický and Helcl, 2017) that can integrate information from different modalities into a coherent output.

MFFG (RNN/Transformer) (Liu et al., 2020). The multistage fusion with forget gate (MFFG) model proposes a cross fusion block with forget gate and a hierarchical fusion decoder to improve multimodal generation.

4.4 Evaluation Metrics

Following (Liu et al., 2020), we use ROUGE, BLEU, METEOR, and CIDEr to evaluate the summaries. ROUGE- $\{1, 2, L\}$ (the standard metrics for abstractive summarization) (Lin and Hovy, 2003) and

BLEU- $\{1, 2, 3, 4\}$ (Papineni et al., 2002) are used to calculate the recall and precision of n-gram overlaps, respectively, between the references and the generated summaries. MENTOR (Denkowski and Lavie, 2011) is used to match the word stems, synonyms and paraphrases between the reference and the generated summary. CIDEr (Vedantam et al., 2015) is an image captioning metric to compute the cosine similarity between TF-IDF weighted n-grams.

In addition, We use Content F1 (Palaskar et al., 2019) to measure the F1 score of the content words of the generated summary based on a monolingual alignment. Firstly, METEOR toolkit (Banerjee and Lavie, 2005; Denkowski and Lavie, 2014) is used to obtain the alignment between the summaries and references. Then, the function words and task-specific stop words are removed from the summaries and references. Finally, the remaining content words from the summaries and references are treated as two bags of words, and the F1 scores are calculated over the alignment. Content F1 focuses more on the content and it can avoid the increase of the ROUGE score from the stop words.

We use `nlg-eval`³ to compute the BLEU, MENTOR and CIDEr scores, and use `rouge`⁴ to compute ROUGE scores. The implementation of Content F1 scores follows (Palaskar et al., 2019).

5 Results and Analysis

5.1 Main Results

From Table 1, we can see that when there is only transcript in the input data, S2S and PG reach similar scores in terms of all evaluation metrics. This could be attributed to the fact that PG tends to copy the content in the transcripts while the reference summaries in the How2 dataset have a great number of novel n-grams, which are defined to be novel with respect to the transcript. We also observe that TF performs better than RNN-based models. It is because TF can learn better relationships between words by multi-head attention mechanism and positional embeddings. Furthermore, both text-only T5 and BART outperform all the baseline models by a large gap owe to their pre-trained text generation ability. Compared to T5, BART achieves higher scores mainly because it introduces a novel pre-training objective named sentence permutation.

³<https://github.com/Maluuba/nlg-eval>

⁴<https://github.com/neural-dialogue-metrics/rouge>

Sentence permutation requires the model to generate the original uncorrupted text from randomly shuffled sentences, which enhances the understanding of long text and benefits the summarization task. Moreover, BART is even better than all previous multimodal models trained on transcript and video.

The visual guidance consistently boosts the performance of T5 and BART by a large step. As shown in Table 2, our best model VG-BART+FG+VTF with the cross-modal multi-head attention surpasses the previous state-of-the-art model (MFFG) by 5.7 ROUGE-1, 5.3 ROUGE-2, and 5.1 ROUGE-L scores. The visual guidance contributes 83.6% of the overall improvement on average of all ROUGE scores.

The results of Content F1 scores in Table 1 show similar trends with other evaluation metrics. By injecting visual information, the models can generate summaries with much richer content. Table 2 shows that both forget gate (FG) and visual transformer encoder (VTF) benefit the model’s performance. However, the Content F1 score is not boosted when combining FG and VTF together, which is contradictory to all other metrics. We conjecture that it is because the Content F1 focuses more on the content aspect, it may have some variance compare to other metrics.

5.2 How to Inject Visual Information

As illustrated in Section 3.3, we mainly adopt two text-vision fusion mechanisms to inject visual information, the cross-modal dot-product attention and multi-head attention. As shown in Table 1, for the VG-BART model, these two fusion mechanisms consistently improve its performance on all metrics by a comparable margin. However, for the VG-T5 model, the cross-modal dot-product attention based fusion does not show any improvement compared to the text-only T5, while the multi-head attention base fusion still increase its performance. We think there are two reasons behind this phenomenon. Firstly, as discussed in Section 5.1, BART leverages the sentence permutation method as its pre-training objective, which increases its robustness on attention-based fusion. Secondly, multi-head attention can capture different key components in the visual information from multiple aspects, which makes it more potent than the dot-product based fusion. Additionally, as mentioned in Section 3.3, we build a variant of the dot-product attention based fusion, which achieves 66.1 ROUGE-1, 49.3 ROUGE-2

Input	Method	R-1	R-2	R-L
Transcript	T5	62.8	45.0	57.5
	BART	64.0	46.4	58.9
Transcript +Noise	VG-T5 (Dot-product)	62.5	43.9	57.0
	VG-T5 (Multi-head)	62.8	44.6	57.4
	VG-BART (Dot-product)	63.9	45.6	58.6
	VG-BART (Multi-head)	63.9	46.5	58.7
Transcript +Video	VG-T5 (Dot-product)	63.0	44.9	57.6
	VG-T5 (Multi-head)	63.3	45.3	58.0
	VG-BART (Dot-product)	66.1	49.3	61.2
	VG-BART (Multi-head)	66.3	49.4	61.4

Table 3: Results of using uniform noise to replace the visual features.

and 61.4 ROUGE-L on VG-BART. This comparable result shows that the variant does not provide further improvement.

To ensure the visual features really help in the learning and our add-on layers aid the understanding of them, we conduct further experiments by replacing the visual features in the input data with random noise of the same dimension and sequence length. The noise is sampled from a uniform distribution from 0 to 3, in a similar value range of the original visual features. As depicted in Table 3, VG GPLMs with random noise as visual features achieve similar or slightly worse performance compared to the text-only GPLMs. This shows the effectiveness of our method to keep GPLMs’ text generation ability. Furthermore, compared to the dot-product attention based fusion, the multi-head fusion is better at retaining GPLMs’ performance, which again demonstrates its superiority.

As mentioned in Section 3, we use a forget gate (FG) to deal with the redundancy and noisy information in the visual features. Additionally, we further encode the visual features by a visual transformer encoder (VTF). Table 2 shows that using either FG or VTF can increase the performance of VG-BART. Jointly leveraging them boosts the performance by 1.7, 2.0, and 1.9 of ROUGE-1, ROUGE-2, and ROUGE-L, respectively.

5.3 Where to Inject Visual Information

As discussed in Section 1, one of the main challenges of building VG GPLMs is to find the optimal location to inject the visual information (i.e., the text-vision fusion). A sub-optimal location might lead to a less effective modality fusion and even hurt the GPLMs’ original text generation ability. As GPLMs have a stack of layers in the encoder

Encoder Layer (BART-base)						R-1	R-2	R-L
1	2	3	4	5	6			
✗	✗	✗	✗	✗	✗	64.0	46.4	58.9
✓	✗	✗	✗	✗	✗	66.7	49.9	61.8
✗	✓	✗	✗	✗	✗	67.0	50.5	62.2
✗	✗	✓	✗	✗	✗	67.3	50.8	62.4
✗	✗	✗	✓	✗	✗	67.4	50.9	62.6
✗	✗	✗	✗	✓	✗	67.4	50.8	62.5
✗	✗	✗	✗	✗	✓	67.7	51.3	63.0
✓	✓	✓	✓	✓	✓	60.4	43.4	55.8
✗	✓	✓	✓	✓	✓	64.1	47.0	59.3
✗	✗	✓	✓	✓	✓	65.3	49.2	60.0
✗	✗	✗	✓	✓	✓	67.5	50.9	62.7
✗	✗	✗	✗	✓	✓	68.0	51.4	63.3

Table 4: Performance of different text-vision fusion locations in the encoder of our best model (VG-BART+FG+VTF with cross-modal multi-head attention). ✓ indicates the occurrence of fusion at a certain layer and ✗ indicates non-occurrence. The first row is the result of BART using transcript only.

and also the decoder, we explore this problem from two aspects: 1) which single layer has the best fusion effect; and 2) does multiple times of fusion help GPLMs to understand the visual information better?

As depicted in Table 4 and 5, firstly, we enumerate each single layer in the encoder and decoder of our best model (VG-BART+FG+VTF) to perform the text-vision fusion. In terms of ROUGE scores, we can clearly tell that injecting visual information into the encoder can generally boost the model’s performance by a large step, while injecting into the decoder only shows negligible improvement. Furthermore, in the encoder, we observe that injecting at a higher layer (closer to the encoder output) brings more improvement. Instead, in the decoder, there is no clear pattern showing the influence of injecting location. We speculate that an early text-vision fusion in the encoder makes the visual information slightly fades away after passing through the stack of encoder layers. Additionally, during the decoding stage, the model utilizes visual information better through the encoder-decoder attention layers than directly injecting into the decoder, which could potentially hurts the generation ability. Secondly, as shown in the lower part of Table 4, we conduct multiple times of fusion in the encoder’s different

Decoder Layer (BART-base)						R-1	R-2	R-L
1	2	3	4	5	6			
×	×	×	×	×	×	64.0	46.4	58.9
✓	×	×	×	×	×	64.6	47.1	59.6
×	✓	×	×	×	×	65.2	48.0	60.3
×	×	✓	×	×	×	64.9	46.9	59.6
×	×	×	✓	×	×	64.8	46.9	59.7
×	×	×	×	✓	×	64.3	46.6	59.1
×	×	×	×	×	✓	64.4	46.7	59.0

Table 5: Performance of different fusion locations in the decoder of our best model (VG-BART+FG+VTF with cross-modal multi-head attention).

locations. We observe that when fusing at all encoder layers simultaneously, the model converges to a much worse performance. We conjecture that this causes the catastrophic forgetting of the pre-trained knowledge in GPLMs. We find that fusing at the last several layers (e.g., 5 and 6) in the encoder is able to further improve the summarization performance.

5.4 Effects of the Forget Gate

As mentioned in Section 3.3, we apply a forget gate (Eq.12) to filter out noise and let the model focus on more important visual information. To have a deeper understanding of the effects of the forget gate, we calculate the average forget gate score (averaged over the whole sequence) for each sample from the How2 test set. As shown in Figure 3, most scores are distributed between 0.47 and 0.48. There is one data sample the score reaches 0.5 because its transcript is not available. As illustrated in Table 6, the model can still generate reasonable summary for it by paying more attention to the visual information. The meaning of the generated summary is still highly aligned with the reference summary, which shows the capability and flexibility of our model to utilize visual information.

6 Conclusion and Future Work

In this paper, we introduce a simple yet effective method to construct vision guided large-scale generative pre-trained language models (VG-BART and VG-T5) for the multimodal abstractive summarization task by inserting attention-based add-on layers. We propose two types of attention mechanisms for the text-vision fusion and interaction: 1) Cross-modal Dot-product Attention; and 2) Cross-modal Multi-head Attention. Moreover, we also

Transcript: transcript not available

Summary from Transcript + Video: learn tips on how to write “cane” in chinese radicals with mandarin characters in the free video clip. get free foreign language lessons from an expert.

Reference Summary: learn what ticks are in chinese calligraphy in this free video clip on languages and writing.

Table 6: An example from How2 testing dataset that has high forget gate score.

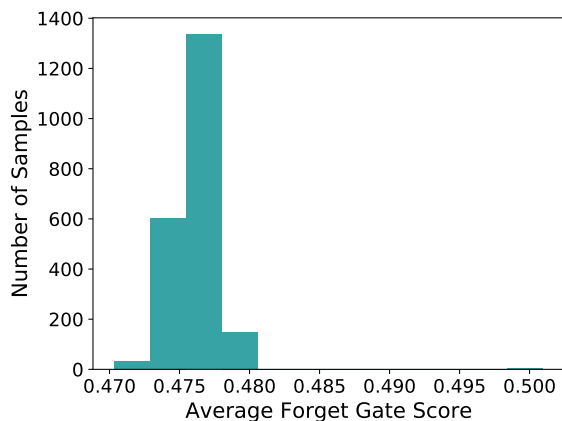


Figure 3: The distribution of average forget gate score on the How2 test set. The model is the VG-BART with dot-product attention.

investigate the effects of using the forget gate and visual transformer encoder along with the attention mechanisms. In addition, we enumerate almost all possible locations in GPLMs for injecting add-on layers. Experimental results show that our approaches significantly outperform the prior state-of-the-art on the How2 dataset. Further analysis illustrates that multi-head attention is more robust than the dot-product attention and higher layers of the encoder is the optimal place to inject vision information. For future work, we believe that our analyses on the *how* and *where* to inject visual information into GPLMs can be applied to other multimodal tasks.

7 Acknowledgments

We want to thank the anonymous reviewers for their constructive feedback. This work is partially funded by ITS/353/19FP and and MRP/055/18 of the Innovation Technology Commission, the Hong Kong SAR Government.

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.
- Jimmy Ba, J. Kiros, and Geoffrey E. Hinton. 2016. Layer normalization. *ArXiv*, abs/1607.06450.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- 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.
- X. Chen, H. Fang, TY Lin, R. Vedantam, S. Gupta, P. Dollár, and C. L. Zitnick. 2015. Microsoft coco captions: Data collection and evaluation server. *arXiv:1504.00325*.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020. Uniter: Universal image-text representation learning. In *ECCV*.
- Jaemin Cho, Jie Lei, Haochen Tan, and M. Bansal. 2021. Unifying vision-and-language tasks via text generation. In *ICML*.
- Wenliang Dai, Samuel Cahyawijaya, Zihan Liu, and Pascale Fung. 2021. **Multimodal end-to-end sparse model for emotion recognition**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5305–5316, Online. Association for Computational Linguistics.
- Wenliang Dai, Zihan Liu, Tiezheng Yu, and Pascale Fung. 2020. **Modality-transferable emotion embeddings for low-resource multimodal emotion recognition**. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 269–280, Suzhou, China. Association for Computational Linguistics.
- Michael Denkowski and Alon Lavie. 2011. Meteor 1.3: Automatic metric for reliable optimization and evaluation of machine translation systems. In *Proceedings of the sixth workshop on statistical machine translation*, pages 85–91.
- Michael Denkowski and Alon Lavie. 2014. Meteor universal: Language specific translation evaluation for any target language. In *Proceedings of the ninth workshop on statistical machine translation*, pages 376–380.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. 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 4171–4186.
- Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. 2015. Long-term recurrent convolutional networks for visual recognition and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2625–2634.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the V in VQA matter: Elevating the role of image understanding in Visual Question Answering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. 2018. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*.
- Dan Hendrycks and Kevin Gimpel. 2016. Gaussian error linear units (gelus). *arXiv: Learning*.
- Zhicheng Huang, Zhaoyang Zeng, Bei Liu, Dongmei Fu, and Jianlong Fu. 2020. Pixel-bert: Aligning image pixels with text by deep multi-modal transformers. *ArXiv*, abs/2004.00849.
- Will Kay, João Carreira, K. Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, T. Back, A. Natsev, Mustafa Suleyman, and Andrew Zisserman. 2017. The kinetics human action video dataset. *ArXiv*, abs/1705.06950.
- Aman Khullar and Udit Arora. 2020. **MAST: Multimodal abstractive summarization with trimodal hierarchical attention**. In *Proceedings of the First International Workshop on Natural Language Processing Beyond Text*, pages 60–69, Online. Association for Computational Linguistics.
- Wonjae Kim, Bokyung Son, and Ildoo Kim. 2021. **Vilt: Vision-and-language transformer without convolution or region supervision**. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 5583–5594. PMLR.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.

- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Gen Li, Nan Duan, Yuejian Fang, Daxin Jiang, and M. Zhou. 2020. Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training. In *AAAI*.
- Haoran Li, Junnan Zhu, Cong Ma, Jiajun Zhang, and Chengqing Zong. 2017. Multi-modal summarization for asynchronous collection of text, image, audio and video. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1092–1102.
- Jindřich Libovický and Jindřich Helcl. 2017. Attention strategies for multi-source sequence-to-sequence learning. *arXiv preprint arXiv:1704.06567*.
- Chin-Yew Lin and Eduard Hovy. 2003. Automatic evaluation of summaries using n-gram co-occurrence statistics. In *Proceedings of the 2003 Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics*, pages 150–157.
- Nayu Liu, Xian Sun, Hongfeng Yu, Wenkai Zhang, and Guangluan Xu. 2020. Multistage fusion with forget gate for multimodal summarization in open-domain videos. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1834–1845.
- 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*.
- Huaishao Luo, Lei Ji, Botian Shi, H. Huang, Nan Duan, Tianrui Li, Xilin Chen, and M. Zhou. 2020. Univilm: A unified video and language pre-training model for multimodal understanding and generation. *ArXiv*, abs/2002.06353.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*.
- Youssef Mroueh, Etienne Marcheret, and Vaibhava Goel. 2015. Deep multimodal learning for audio-visual speech recognition. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2130–2134. IEEE.
- Shruti Palaskar, Jindřich Libovický, Spandana Gella, and Florian Metze. 2019. Multimodal abstractive summarization for how2 videos. *arXiv preprint arXiv:1906.07901*.
- 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.
- Adam Paszke, S. Gross, Francisco Massa, A. Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Z. Lin, N. Gimelshein, L. Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2017. A deep reinforced model for abstractive summarization. *arXiv preprint arXiv:1705.04304*.
- Weizhen Qi, Yu Yan, Yeyun Gong, Dayiheng Liu, Nan Duan, Jiusheng Chen, Ruofei Zhang, and Ming Zhou. 2020. ProphetNet: Predicting future n-gram for sequence-to-SequencePre-training. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2401–2410, Online. Association for Computational Linguistics.
- 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):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Wasifur Rahman, Md Kamrul Hasan, Sangwu Lee, AmirAli Bagher Zadeh, Chengfeng Mao, Louis-Philippe Morency, and Ehsan Hoque. 2020. Integrating multimodal information in large pretrained transformers. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2359–2369, Online. Association for Computational Linguistics.
- Ramon Sanabria, Ozan Caglayan, Shruti Palaskar, Desmond Elliott, Loïc Barrault, Lucia Specia, and Florian Metze. 2018. How2: a large-scale dataset for multimodal language understanding. *arXiv preprint arXiv:1811.00347*.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. *arXiv preprint arXiv:1704.04368*.
- Nitish Srivastava, Geoffrey E. Hinton, A. Krizhevsky, Ilya Sutskever, and R. Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15:1929–1958.

- Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. 2020. [Vi-bert: Pre-training of generic visual-linguistic representations](#). In *International Conference on Learning Representations*.
- Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. 2019. [A corpus for reasoning about natural language grounded in photographs](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6418–6428, Florence, Italy. Association for Computational Linguistics.
- C. Sun, Austin Myers, Carl Vondrick, K. Murphy, and C. Schmid. 2019. Videobert: A joint model for video and language representation learning. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7463–7472.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- 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 5100–5111, Hong Kong, China. Association for Computational Linguistics.
- Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, J. Zico Kolter, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2019. Multimodal transformer for unaligned multimodal language sequences. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Florence, Italy. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017a. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017b. [Attention is all you need](#). *arXiv preprint arXiv:1706.03762*.
- 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.
- Qiang Wang, Bei Li, Tong Xiao, Jingbo Zhu, Changliang Li, Derek F. Wong, and Lidia S. Chao. 2019. [Learning deep transformer models for machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1810–1822, Florence, Italy. Association for Computational Linguistics.
- Wen Xiao and Giuseppe Carenini. 2019. Extractive summarization of long documents by combining global and local context. 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 3002–3012.
- Tiezheng Yu, Zihan Liu, and Pascale Fung. 2021. Adaptsum: Towards low-resource domain adaptation for abstractive summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5892–5904.
- Zhou Yu, Jun Yu, Yuhao Cui, Dacheng Tao, and Qi Tian. 2019. Deep modular co-attention networks for visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Amir Zadeh, Minghai Chen, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. 2017. Tensor fusion network for multimodal sentiment analysis. *arXiv preprint arXiv:1707.07250*.
- Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. From recognition to cognition: Visual commonsense reasoning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020a. [PEGASUS: Pre-training with extracted gap-sentences for abstractive summarization](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 11328–11339. PMLR.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020b. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International Conference on Machine Learning*, pages 11328–11339. PMLR.
- Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. 2021. [Vinvl: Making visual representations matter in vision-language models](#). *CVPR 2021*.
- Luwei Zhou, H. Palangi, Lei Zhang, Houdong Hu, Jason J. Corso, and Jianfeng Gao. 2020. Unified vision-language pre-training for image captioning and vqa. In *AAAI*.

A Multi-head Self-Attention

The query (Q), key (K), value (V) based self-attention is the core building block of the Transformer model (Vaswani et al., 2017b). Given the

input $Z \in \mathbb{R}^{s \times d_z}$, we calculate Q , K , and V by

$$\begin{aligned} Q &= ZW_q, Q \in \mathbb{R}^{s \times d_k} \\ K &= ZW_k, K \in \mathbb{R}^{s \times d_k} \\ V &= ZW_v, V \in \mathbb{R}^{s \times d_v}, \end{aligned}$$

in which $W_q \in \mathbb{R}^{d_z \times d_k}$, $W_k \in \mathbb{R}^{d_z \times d_k}$, and $W_v \in \mathbb{R}^{d_z \times d_v}$ are the projection weights. Then, a single-head self-attention is calculated by

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V,$$

where $\frac{1}{\sqrt{d_k}}$ is the scaling factor to mitigate the extremely small gradients issue mentioned by [Vaswani et al. \(2017b\)](#). For multi-head self-attention, it can be calculated by

$$\text{MultiHead}(Q, K, V) = \text{Concat}(\text{head}_1, \dots, \text{head}_h)W_o$$

and

$$\text{head}_i = \text{Attention}(QW_q^i, KW_k^i, VW_v^i).$$

B Feed-Forward Network

Given the input $Z \in \mathbb{R}^{s \times d_z}$, the feed-forward network (FFN) processes it with two linear projections $W_f^1 \in \mathbb{R}^{d_z \times d_f}$, $W_f^2 \in \mathbb{R}^{d_f \times d_z}$ and a non-linear function GELUs ([Hendrycks and Gimpel, 2016](#)),

$$\text{FFN}(Z) = \text{GELU}(ZW_f^1)W_f^2.$$

In addition, after each linear projection, there is a dropout ([Srivastava et al., 2014](#)) layer to improve generalization.