

Visualize Before You Write: Imagination-Guided Open-Ended Text Generation

Wanrong Zhu[¶], An Yan[†], Yujie Lu[¶], Wenda Xu[¶],
Xin Eric Wang[§], Miguel Eckstein[¶], William Yang Wang[¶]

[¶]UC Santa Barbara, [†]UC San Diego, [§]UC Santa Cruz

{wanrongzhu,yujielu,wendaxu,william}@cs.ucsb.edu, ayan@ucsd.edu

xwang366@ucsc.edu, miguel.eckstein@psych.ucsb.edu

Abstract

Recent advances in text-to-image synthesis make it possible to visualize machine imaginations for a given context. On the other hand, when generating text, human writers are gifted at creative visualization, which enhances their writings by forming imaginations as blueprints before putting down the stories in words. Inspired by such a cognitive process, we ask the natural question of whether we can endow machines with the same ability to utilize visual information and construct a general picture of the context to guide text generation. In this work, we propose iNLG that uses machine-generated images to guide language models (LM) in open-ended text generation. The experiments and analyses demonstrate the effectiveness of iNLG on open-ended text generation tasks, including text completion, story generation, and concept-to-text generation in both few-shot and full-data scenarios. Both automatic metrics and human evaluations verify that the text snippets generated by our iNLG are coherent and informative while displaying minor degeneration.¹

1 Introduction

One great resource human writers cherish is the ability of imagination, with which they render mental images about an actual or vicarious experience and link knowledge that would later make the writing more concrete, sensible, and intriguing. Cognitive studies show that visual imagery improves comprehension during language processing (Gambrell and Bales, 1986; Joffe et al., 2007; Sadoski and Paivio, 2000), and that mental imagery facilitates humans’ written language expression at young ages (Gambrell and Koskinen, 2002).

When it comes to the study of Artificial Intelligence (AI), one classic challenge for AI systems is to generate informative and coherent text snippets. Open-ended text generation is such a task that provides an input context, and asks the model to

¹Our code & data: <https://github.com/VegB/iNLG>.

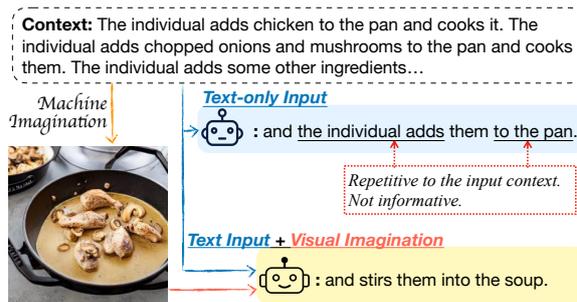


Figure 1: When performing open-ended text generation, the language models prompted with text-only input may generate repetitive or unilluminating contents, which is also known as degeneration. Hereby, we propose to use machine-generated images as additional visual supervision to guide the language models in generating more informative and coherent text with the given context.

generate a piece of text that is consistent with the context. This is the cornerstone of a wide range of downstream tasks such as text completion (Guan et al., 2019; Radford et al., 2019), story generation (Fan et al., 2018; Goldfarb-Tarrant et al., 2020; Swanson et al., 2021; Su et al., 2022b), and dialogue systems (Schatzmann et al., 2007; Wen et al., 2015, 2017; Wei et al., 2018; Wu et al., 2021), and has received much attention throughout the years. Inspired by human writers’ common practice of creative visualization, we ask the following question: Can we endow machines with the same ability to construct a general picture of the context and use it as a blueprint to guide text generation?

Recent advances in text-to-image generation make it possible to visualize machine imaginations for a given context (Ramesh et al., 2021; Rombach et al., 2022; Crowson et al., 2022; Wang et al., 2022b; Saharia et al., 2022). Moreover, this line of work shows great potential in utilizing textual information to guide image synthesis. It comes naturally that one may attempt to complete the loop by using visual supervision to guide text generation.

In this work, we propose using machine-

generated images to guide the language model (LM) in open-ended text generation. More specifically, we visualize machine imagination for the input context by rendering images with StableDiffusion (Rombach et al., 2022), a state-of-the-art text-to-image generator. The machine imagination acts as additional visual supervision to guide LMs in generating informative and coherent text in two ways. Firstly, the machine-generated images are introduced as the input to the LM in the form of the visual prefix. Secondly, we designed a contrastive training objective that enforces the generated text to be semantically similar to the visual supervision.

We conduct experiments on three open-ended text generation tasks, namely text completion, story generation, and concept-to-text generation. Extensive experiments in the few-shot settings show better or competitive performance to state-of-the-art baselines on both automatic metrics and human evaluation. Experiments with full-data settings show that introducing machine-generated visual supervision with our iNLG yields consistent improvements on various LM models including GPT-2 (Radford et al., 2019), BART (Lewis et al., 2020), and T5 (Raffel et al., 2020).

Our main contributions are as follows:

- We introduce a novel paradigm that leverages machine-generated images to guide open-ended text generation. This endows the machines with the ability of creative visualization that human writers often demonstrate.
- We distill the vision information from the pre-trained multimodal models and further construct visual prefixes to guide language models performing text generation with teacher forcing and contrastive objectives.
- Extensive experiments show the effectiveness of iNLG as a model-agnostic framework in open-ended text generation tasks, including text completion, story generation, and concept-to-text in both few-shot and full-data settings.

2 Related Work

Open-ended Conditional Text Generation is the task of generating a coherent portion of the text based on the given context. Recent advances in pre-trained models have pushed frontier in the open-ended conditional text generation, such as text completion (See et al., 2019; Ippolito et al., 2020), story generation (Guan et al., 2020; Fan

et al., 2018; Yao et al., 2019) and concept-to-text generation (Zhou et al., 2021; Liu et al., 2021). Despite the success of large language models, text degeneration and semantic coverage still remain as two core technical challenges in few-shot open-ended text generation. To improve the text coverage, StoryEndGen (Guan et al., 2019) leverages the knowledge graph to encode context sequentially. Fan et al. (2018) and Yao et al. (2019) plan the content (premise or keywords) first and then encourage the generation based on planned content. To mitigate the text degeneration, SimCTG (Su et al., 2022b) uses a contrastive training strategy to encourage the model to learn isotropic token embeddings. Similar to our approach, Wang et al. (2022a) generates a scene graph for each concept and combines them with text for the model input. Previous work has proposed to add visual information to LM by retrieving images from the Internet or large-scale image sets (Yang et al., 2020; Cho et al., 2021; Su et al., 2022a). However, the retrieved images may fail to fully incorporate the context, which will misguide the LM from yielding contextually consistent predictions.² Unlike prior work, our approach leverages images generated conditioning on the context to assist the text generation process.

Visually-aided NLP Recent work show the power of visual guidance in natural language processing, spanning from the language representation learning (Lu et al., 2019; Li et al., 2019; Sun et al., 2019; Luo et al., 2020; Chen et al., 2020; Li et al., 2020; Tan and Bansal, 2020; Lu et al., 2022), the downstream tasks (Grubinger et al., 2006; Elliott et al., 2016; Xie et al., 2019; Christie et al., 2016; Shi et al., 2019; Lu et al., 2022) and evaluation (Zhu et al., 2021). They either leverage visual information from an external vision-and-language corpus or obtain such visual knowledge from the large pre-trained model. In this line of work, imagination achieves promising performance in various NLP domains (Long et al., 2021; Zhu et al., 2021; Wang et al., 2022a; Lu et al., 2022). Previous imagination-based work in NLP either study non-generation problems (Zhu et al., 2021; Lu et al., 2022) or utilize non-visual information (Long et al., 2021; Wang et al., 2022a). Our work explores the potential of generating visual imagination to improve open-ended text generation tasks.

²Figure 8 shows examples where the image retrieved from the search engine is irrelevant with the input context.

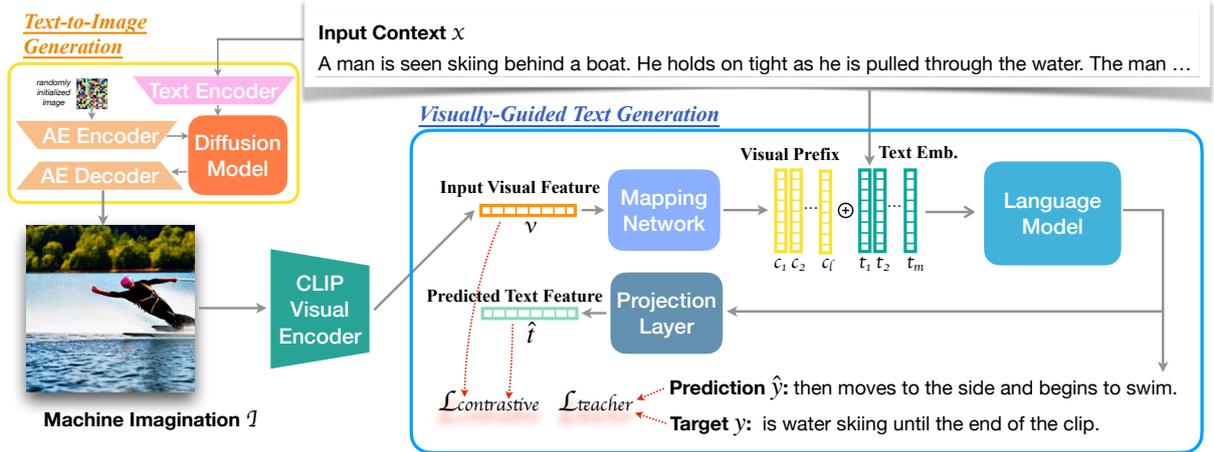


Figure 2: An overview of our iNLG. Given an input context x , we first visualize the context with the text-to-image generation model. Then we use the machine-generated image I as the additional visual supervision to guide the language model in open-ended text generation. The visual feature is provided as a source of input to the LM in the form of the visual prefix. Aside from the teacher forcing objective $\mathcal{L}_{\text{teacher}}$, we also enforce the LM to generate text that is semantically similar to the machine imagination with a contrastive training objective $\mathcal{L}_{\text{contrastive}}$.

3 Method

3.1 Overview

Open-ended text generation is a task that provides an input context, and asks the model to generate a piece of text that is consistent with the context.

This work mainly focused on introducing machine-rendered images to assist LM in performing open-ended text generation. More specifically, given the context x^i , we first use a text-to-image generator to illustrate an image I^i that depicts the input context. The LM is prompted with image I^i as the visual prefix along with the text context x^i , and will incorporate the multimodal input to generate the output text \hat{y}^i .

Figure 2 provides an overview of our iNLG framework, which mainly involves two modules. The first module is a text-to-image generator that takes in the input context and illustrates a descriptive image, which we also refer to as the machine imagination. The second module is a visually-guided language model that utilizes the machine imagination as a source of input and also a supervision that encourages the LM to generate text that is semantically similar to the visual information.

3.2 Text-to-Image Rendering

In this work, we propose to use images generated conditioning on the context by the machines as additional visual information to the LM. The text-to-image generation backbone is StableDiffusion (Rombach et al., 2022), which mainly consists

of a text encoder, a diffusion model, and an auto-encoder. The text encoder is from the frozen CLIP ViT-L/14 (Radford et al., 2021) and encodes the input text to textual embeddings. The diffusion model uses UNet (Ronneberger et al., 2015) to provide noise estimation. The UNet is modified so as to attend to the input textual embeddings. The encoder of the pretrained autoencoder encodes images into the lower-resolution latent maps z_T . At each step t , the diffusion model provides the noise estimation ϵ and modifies z_t correspondingly. The decoder of the pretrained autoencoder takes the final noise-free latent map z and generates the image prediction. StableDiffusion is trained with LAION-5B (Schuhmann et al., 2022).

3.3 Visually Guided Text Generation

Visual Prefix Construction One can encode the visual information with the pre-trained visual models. However, such visual embedding may lie in a representation space different from the LM due to the discrepancy between models. One way of introducing features extracted by another network to the current model is through feature mapping (Mokady et al., 2021). With a dataset of image-text pairs (I', x') , we can pre-train a mapping network \mathcal{F} for a given LM in an image captioning formulation. More specifically, we encode I' with the visual encoder $\text{Enc}_{\text{visual}}$ and receive its visual features v' . Then we apply the mapping network \mathcal{F} over v' , and receive a sequence of l visual prefixes:

$$c'_1, c'_2, \dots, c'_l = \mathcal{F}(v') = \mathcal{F}(\text{Enc}_{\text{visual}}(I')) \quad (1)$$

We provide the list of visual prefix as input to the LM with the corresponding text x' as the target output. Such a pre-training process enables \mathcal{F} to project visual features into the visual prefix that lies within the same embedding distributions as the LM. The mapping network is agnostic of the downstream task, and only depends on the visual source and the LM.

After generating a descriptive image I^i for the input context x^i , we use CLIP to encode I^i and receive its visual features v^i . We apply the pre-trained mapping network \mathcal{F} over v^i , and receive the visual prefix c^i of length l :

$$c^i = \{c_1^i, c_2^i, \dots, c_l^i\} = \mathcal{F}(\text{CLIP}(I^i)) \quad (2)$$

Visually-guided Language Modeling We use the visual information to guide text generation in two ways, reflected in the following two training objectives. Firstly, we directly introduce the machine-generated visual information as input to the LM. We concatenate the visual prefix c^i and the text embeddings t^i for the input context x^i with m tokens. LM input can be denoted as $[c^i; t^i] = \{c_1^i, \dots, c_l^i, t_1^i, \dots, t_m^i\}$. With $y^i = \{y_1^i, y_2^i, \dots, y_n^i\}$ denoting the target output of n tokens, and θ denoting the trainable parameters, we can list out the teacher forcing training objective as follows:

$$\mathcal{L}_{\text{teacher}} = - \sum_{j=1}^n \log p_{\theta}(y_j^i | c^i; t^i; \mathbf{y}_{<j}^i) \quad (3)$$

In addition, we design a contrastive objective to enforce the generated text to be semantically similar to the input visual supervision with the InfoNCE loss (van den Oord et al., 2018; Yan et al., 2021):

$$\mathcal{L}_{\text{contrastive}} = - \log \frac{\exp(\text{sim}(\mathbf{v}^i, \hat{\mathbf{t}}^i)/\tau)}{\sum_{j \neq i} \exp(\text{sim}(\mathbf{v}^i, \hat{\mathbf{t}}^j)/\tau)} \quad (4)$$

in which $\hat{\mathbf{t}}$ is the projected representation of the decoder’s last layer’s output, and can be viewed as the sentence-level representation of the generated text. Here $\text{sim}(\cdot, \cdot)$ first normalizes the two vectors, then compute their cosine similarity, and τ is the temperature.

3.4 Training & Inference

We first pre-train the mapping network on the pre-training dataset with the teacher-forcing objective. Such pre-training is agnostic of the downstream task, and only depends on the type of base LM.

When applying our iNLG on downstream tasks, we train the base LM with the teacher forcing objective for the first $N_{\text{no_contra}}$ epochs. Then, we introduce the contrastive objective and tune the base LM together with the mapping network and projection layer by minimizing the following loss \mathcal{L} . Here ep denotes the epoch and λ is the factor:

$$\mathcal{L} = \begin{cases} \mathcal{L}_{\text{teacher}}, & ep < N_{\text{no_contra}}, \\ \mathcal{L}_{\text{teacher}} + \lambda \mathcal{L}_{\text{contrastive}}, & ep > N_{\text{no_contra}}, \end{cases} \quad (5)$$

During inference, we provide the context and machine-generated image to the LM. We use beam search during decoding with a beam width of 10.

4 Experimental Setup

4.1 Tasks, Datasets, and Baselines

We apply our iNLG on three open-ended text generation setups: sentence completion, story generation, and concept-to-text generation. Table 1 shows examples for each task.

Sentence Completion is a task of finishing the sentence in a commonsense inference scenario. We conduct experiments on the ActivityNet (Heilbron et al., 2015) subset³ of HellaSwag (Zellers et al., 2019), which is a benchmark for commonsense natural language inference that ask the model to predict the most likely follow-up among several choices given a specific context. We compare with StoryEndGen (Guan et al., 2019) which encodes the given context incrementally and attends to the one-hop knowledge graph retrieved from ConceptNet for the context tokens. We implement our iNLG on top of the GPT-2 (Radford et al., 2019), which by nature, can generate the follow-up for an arbitrary input in a zero-shot manner.

Story Generation requires the model to compose a story based on the given title or context. We conduct experiments on the widely used story generation benchmark ROCStories (Mostafazadeh et al., 2016). Each data item consists of a story title and a human-written five-sentence everyday life story that incorporates commonsense related to the title.⁴ We provide the story title and the story’s first sentence as the input context, and ask the LM to predict the following four sentences. We consider the

³14740/982/2261 samples for train/validation/test.

⁴We use the split provided by Su et al. (2022a), which is based on the ROCStories Winter 2017 release and contains 49666/1500/1500 items for the train/validation/test sets.

Task	Input Context	Target Output
Text Completion	Different people are interviewed on camera while several others are shown raking up the leaves. A man is seen sitting in his car and another puts his gloves on. The camera	pans over the raked up leaves while several others discuss their hard work.
Story Generation	Live Show. Tim was in his school’s play.	He was nervous about their first show. He almost dropped out. The show went smoothly. Tim was excited for his second show.
Concept-to-Text	grow, flower, pavement	Wild flower growing through crack in the tiled pavement.

Table 1: Exemplars of the input context and corresponding target output for three open-ended text generation task covered in this study, namely story generation, text completion, and concept-to-text generation.

following methods as baselines: Action-Plan (Fan et al., 2018) first predicts the premise of a story with the convolutional LM (Dauphin et al., 2017), then use the fusion mechanism (Sriram et al., 2018) to encourage a convolutional seq2seq model (Gehring et al., 2017) to generate the story from the premise. Plan-and-Write (Yao et al., 2019) first plans a storyline that consists of keywords, then generate the story conditioned on the storyline. Its model structure is built upon GRU (Cho et al., 2014). SimCTG (Su et al., 2022b) proposes a contrastive training objective that encourages the LM to learn discriminative and isotropic token representations, and is implemented on GPT-2 (Radford et al., 2019).

Concept-to-Text is a relatively more constrained conditional text generation task involving common-sense reasoning. This task provides a set of concepts as input, and requires the model to generate a piece of text that incorporates the concepts and describes an everyday scenario. We conduct experiments on the CommonGen (Lin et al., 2020) benchmark.⁵ We compare against the following models: KG-BART (Liu et al., 2021) encompasses the relations of concepts with the knowledge graph and augments the BART (Lewis et al., 2020) encoder and decoder with graph representations. ModelAdapt (Ma et al., 2021) is built upon BART and removes the positional embedding in the encoder. Imagine-and-Verbalize (I&V) (Wang et al., 2022a) predicts a scene graph for each set of concepts, and uses it as an additional input to the LM. In contrast to I&V, we directly visualize the concepts and use the machine-generated images as the auxiliary information to assist the concept-to-text generation.

4.2 Evaluation

Automatic For sentence completion and story generation, we follow previous work and eval-

⁵We use the in-house split provided by Wang et al. (2022a), which contains 65323/2066/4018 samples for train/validation/test.

uate the quality of the generated text from the aspect of model degeneration level (rep- n , diversity, distinct- n), text distribution divergence (MAUVE), and semantic similarity (BERTScore): (1) rep- $n = 1.0 - \frac{|\text{unique } n\text{-grams}|}{|\text{total } n\text{-grams}|}$ measures sequence level repetition by computing the portion of duplicate n -grams (Welleck et al., 2020). (2) diversity = $\prod_{n=2}^4 (1 - \text{rep-}n)$ measures the diversity of n -grams (Su et al., 2022a). (3) distinct- $n = \frac{|\text{unique } n\text{-grams}|}{|\text{length of text}|}$ measures the portion of distinct n -grams in the text (Li et al., 2016). (4) MAUVE measures the learned distributions divergence between the generated text and human-written text (Pillutla et al., 2021),⁶ a low MAUVE indicates a great difference between the distributions of generated text and human text. (5) BERTScore assesses contextual text similarity between two pieces of texts by computing the cosine similarities between their tokens’ embeddings (Zhang* et al., 2020),⁷ a low BERTScore means the generated text is contextually different from the ground-truth.

For concept-to-text, following prior work, we report the metrics scores on BLEU (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), CIDEr (Vedantam et al., 2015), SPICE (Anderson et al., 2016), and BERTScore (Zhang* et al., 2020).

Human We also set up a human evaluation as a complementary evaluation beyond the automatic metrics. We select 100 samples from the test set for sentence completion and story generation and perform the head-to-head comparison between the text snippets generated by our iNLG and the baseline models. We invite human annotators to compare the text quality from the following three independent aspects: (1) *Coherence*: Which snippet is more semantically consistent with the context, and follows the logic of the context more naturally. (2) *Fluency*: Which snippet is more fluent in English. (3) *Informativeness*: Which snippet contains more

⁶We report MAUVE with gpt2-large as the base model.

⁷We report BERTScore with roberta-large as base model.

Task	*	Setting	rep-2 ↓	rep-3 ↓	rep-4 ↓	diversity ↑	distinct-2 ↑	MAUVE ↑	BERTScore ↑
Sentence Completion	0	Human	0.45	0.05	0.01	99.50	77.32	-	-
	1	GPT2 <i>no finetune</i> (Radford et al., 2019)	6.71	6.87	10.13	78.07	74.83	44.19	22.57
	2	StoryEndGen (Guan et al., 2019)	39.53	35.11	39.30	34.12	44.57	0.45	-47.29
	3	GPT2 <i>text-only finetune</i>	4.20	4.03	5.53	86.85	75.14	49.45	24.13
	4	GPT2 +iNLG	2.43	2.61	3.57	91.63	75.92	60.30	24.25
Story Generation	5	Human	1.76	0.38	0.15	97.71	56.34	-	-
	6	GPT2 <i>no finetune</i>	37.65	22.76	21.92	45.67	43.42	0.43	-7.77
	7	Action-Plan (Fan et al., 2018)	52.05	35.58	28.11	26.97	21.43	0.41	-18.32
	8	Plan-and-Write (Yao et al., 2019)	45.22	32.86	23.34	30.71	20.83	0.41	-37.35
	9	SimCTG (Su et al., 2022b)	28.72	24.02	20.61	43.00	42.06	0.43	18.01
	10	GPT2 <i>text-only finetune</i>	25.41	18.51	14.41	52.10	46.60	9.10	21.23
	11	GPT2 +iNLG	10.73	5.64	3.42	81.36	51.91	35.94	23.03

Table 2: Generation quality scores for few-shot text completion on the ActivityNet and few-shot story generation on ROCStories. “Human” shows the human performance and “GPT2 *no finetune*” denotes the vanilla GPT2 model without tuning. All the other listed models are trained with 1% of the training data. “+iNLG” denotes introducing machine-generated images on top of the base LM.

interesting content, and describes the scenes that are more likely to happen in real life. Three human judges rate each comparison.

4.3 Implementation Details

We use StableDiffusion-v1-1 (Rombach et al., 2022) to render a 512x512 image from the context, and use CLIP ViT/B-32 to extract features offline. The mapping network is an 8-layer Transformer, and the visual prefix length is 20. For the sentence completion and story generation tasks, the mapping network is pre-trained on the MSCOCO (Lin et al., 2014) dataset. For the concept-to-text task, the mapping network is pre-trained on VIST (Huang et al., 2016).⁸ We pre-train the mapping network for 5 epochs with a batch size of 128. Results are reported on three repeat runs. Detailed hyperparameters are listed in the Appendix.

5 Result and Analysis

5.1 Few-Shot Learning Results

Open-ended text generation is a broad topic with flexible and inexhaustible setups, many of which have low resources. Collecting annotations is often extremely expensive and time-consuming. Therefore, we first report few-shot results to check if our iNLG can rapidly adapt to new task setups with a few examples, which is more practical in real-life.

More specifically, we report few-shot open-ended text generation results with 1% of the training data. For sentence completion and story gen-

eration tasks, the base LM is GPT2-base (Radford et al., 2019). For concept-to-text, we test it with BART-base (Lewis et al., 2020) as the base LM.

Sentence Completion As shown in Table 2, StoryEndGen (#2) suffers from degeneration with the highest rep- n and the lowest diversity. Training with only 1% of the training data improves GPT2’s performance on all metrics (#3 vs. #1). Under the same few-shot setting, adding additional machine-generated images with our iNLG (#4) further alleviate model degeneration. The improvement on MAUVE also indicates that introducing visual input can aid GPT2 in generating text that is more similar to the human-written ones.

Story Generation As shown in Table 2, for the story generation task that requires the LM to compose longer text, we see the vanilla GPT2 without tuning suffering from more severe degeneration compared to rendering a sentence ending (#6 vs. #1). The high rep- n scores indicate that the two non-Transformer-based baselines Action-Plan (#7) and Plan-and-Write (#8) stammer with repetitive tokens, which greatly differs from the human-written text (leads to low MAUVE) and does not have concrete meanings (leads to low BERTScore). The models based on GPT-2 (#9-#10) yield more complete sentences with concrete meanings (BERTScore gets higher). However, they keep repeating the same sentence, which is still quite different from human language (MAUVE remains low). Applying iNLG to GPT-2 leads to minor degeneration and has the best performance on all metrics (#11). Examples of generated text snippets can be found in Figure 6 and in Appendix.

⁸CommonGen is built upon image and video captioning datasets including MSCOCO. To avoid data leakage, we choose to pre-train the mapping network on VIST, which is not revealed to CommonGen.

Task	Models	Coherence			Fluency			Informativeness		
		Win(%)	Tie(%)	Lose(%)	Win(%)	Tie(%)	Lose(%)	Win(%)	Tie(%)	Lose(%)
Sentence Completion	Ours vs. StoryEndGen	51.67	20.33	28.00	44.67	19.33	36.00	41.33	18.33	40.33
	Ours vs. GPT2 <i>no finetune</i>	51.00	22.67	26.33	45.00	22.33	32.67	41.00	21.00	38.00
	Ours vs. GPT2 <i>text-only finetune</i>	58.00	24.33	17.67	43.33	18.67	38.00	42.33	21.67	36.00
Story Generation	Ours vs. Action-Plan	51.00	24.67	24.33	54.67	16.33	29.00	52.00	15.00	33.00
	Ours vs. Plan-and-Write	45.33	25.67	29.00	53.00	16.67	30.33	54.67	17.00	28.33
	Ours vs. SimCTG	42.00	27.67	30.33	40.33	25.67	34.00	43.33	18.33	38.33
	Ours vs. GPT2 <i>no finetune</i>	43.33	24.33	32.33	43.67	20.33	36.00	44.67	19.00	36.33
	Ours vs. GPT2 <i>text-only finetune</i>	39.33	26.67	34.00	38.67	26.67	34.67	44.33	22.67	33.00

Table 3: Human evaluation results for the sentence completion task and the story generation task. The scores indicate the percentage of win, tie or lose when comparing our iNLG with the baseline models.

* Setting	B-4	M.	CIDEr	SPICE	BertS.
1 BART-base <i>text-only finetune</i>	20.72	25.47	114.49	24.58	59.76
2 +KG (Liu et al., 2021)	15.26	24.44	98.53	23.13	52.76
3 +Adapt (Ma et al., 2021)	23.11	25.96	123.44	25.14	61.53
4 +I&V (Wang et al., 2022a)	24.50	25.89	119.61	25.59	57.29
5 +iNLG	25.07	26.48	127.93	26.32	63.37

Table 4: Automatic metrics scores for few-shot concept-to-text generation on CommonGen with 1% of the training data. All listed models are implemented on BART-base. “+KG” adds knowledge graph, “+Adapt” applies model adaption, “+I&V” adds scene graph, and “+iNLG” introduces machine-generated images as input. B-4: BLEU-4; M.: METEOR; BertS.: BERTScore.

Concept-to-Text Table 4 shows that knowledge graph information may not be fully exploited under the few-shot setting (#2), while removing the information of relative positions between input concepts helps the LM write better sentences (#3). Introducing machine-generated images can improve the base LM’s performance on concept-to-text generation (#5 vs. #1). While both I&V and our iNLG involve machine “imagination”, we provide such information in different forms (scene graphs vs. images). Comparing #4 and #5, our iNLG outperforms I&V with BART-base as the base LM. This suggests that the additional information introduced by I&V and iNLG is complementary.

Human Evaluation Table 3 lists out human evaluation results on text completion and story generation. Our iNLG outperforms the compared baselines on all three criteria in the model-level head-to-head comparisons. This further verifies the effectiveness of our iNLG in generating fluent and informative text snippets that better align with the given context.

5.2 Model-Agnostic Improvement

We further report open-ended text generation results with various base LM when trained with the full set of data. For concept-to-text, we experiment

Base LM	Setting	Metrics				
<i>Concept-to-Text</i>		B-4↑	MET↑	CIDEr↑	SPICE↑	BertS.↑
BART-base	text-only	30.32	31.35	158.92	31.22	68.50
	+iNLG	30.60	31.44	160.63	31.42	69.02
BART-large	text-only	32.38	33.06	169.69	33.01	70.33
	+iNLG	32.76	33.17	171.47	33.35	70.79
T5-base	text-only	30.39	30.87	163.67	32.77	70.03
	+iNLG	31.09	31.18	165.52	32.81	70.35
T5-large	text-only	34.13	32.91	175.67	34.30	72.44
	+iNLG	34.50	33.87	177.65	35.48	72.70
<i>Sentence Completion</i>		rep-4↓	div.↑	dist-2↑	MAUVE↑	BertS.↑
GPT2-base	text-only	4.20	87.46	72.87	61.42	29.84
	+iNLG	3.95	89.33	74.09	64.01	30.10
GPT2-large	text-only	1.77	96.54	76.74	87.81	31.66
	+iNLG	2.05	95.90	76.80	89.11	32.15
<i>Story Generation</i>		rep-4↓	div.↑	dist-2↑	MAUVE↑	BertS.↑
GPT2-base	text-only	7.83	68.42	49.53	33.13	28.81
	+iNLG	6.80	71.17	49.92	38.86	29.13
GPT2-large	text-only	1.02	91.91	54.17	82.81	31.86
	+iNLG	0.85	92.51	54.54	87.83	32.03

Table 5: Automatic metric scores when trained with the full set of data with ablations of the base LM. Introducing our iNLG leads to model-agnostic improvements across the board. B-4: BLEU-4; MET.: METEOR; BertS.: BERTScore; div.: diversity; dist-2: distinct-2.

with BART-base/large (Lewis et al., 2020) and T5-base/large (Raffel et al., 2020). For sentence completion and story generation, we record results on GPT2-base/large (Radford et al., 2019). As shown in Table 5, introducing machine-generated visual supervision with our iNLG leads to model-agnostic improvements over text-only finetuning. This holds true for all the listed base LM with different architectures and verifies that our iNLG is a model-agnostic framework.

5.3 Performance Analysis

Source of Image We first perform an ablation study to understand how the source of visual information affects our iNLG framework. We compare retrieved/generated images from four sources: (1) the first returned result by Yahoo Image Search;⁹

⁹<https://images.search.yahoo.com/>

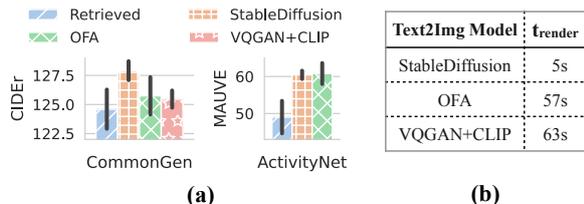


Figure 3: (a) iNLG’s performance on CommonGen and ActivityNet with visual supervisions retrieved from the web or generated by machines. Scores are reported with error bars. (b) Average time to render an image on TITAN RTX with each text-to-image generator.

(2) images rendered by VQGAN+CLIP (Crowson et al., 2022);¹⁰ (3) images rendered by OFA (Wang et al., 2022b),¹¹ and (4) images rendered by StableDiffusion (Rombach et al., 2022), with which we report the main results.

As shown in Figure 3(a), the images generated by machines act as a more effective supervision than the retrieved images. This validates our motivation of introducing machine-generated images over retrieved ones to guide LM in performing text generation. Among the three text-to-image generators, VQGAN+CLIP is slightly inferior to the other two, while StableDiffusion and OFA have mixed performance. Images generated by StableDiffusion rank first on CommonGen, while images rendered with OFA score slightly higher on ActivityNet. Figure 3(b) reports the average image rendering time, where StableDiffusion is 10× faster when rendering images than the other two.

Contrastive Training We examine the effect of the contrastive training objective on CommonGen, and the results are presented in Figure 4. We notice that introducing $\mathcal{L}_{\text{contrastive}}$ improves iNLG’s performance on 4 out of 5 listed few-shot setups, which suggests that our contrastive training objective generally can assist the LM in composing open-ended text snippets. One exception is in the extreme few-shot setting with only 0.1% of training data, where the amount of data is insufficient to let the LM form a decent representation. In this case, enforcing the sentence representation to be similar to the visual supervision with $\mathcal{L}_{\text{contrastive}}$ might misguide the LM.

Mapping Network & Visual Prefix We discuss the effects of different types of mapping networks and various visual prefix lengths. Aside from the 8-layer Transformer we used in the main experi-

¹⁰<https://github.com/nerdyrodent/VQGAN-CLIP>

¹¹<https://github.com/OFA-Sys/OFA>

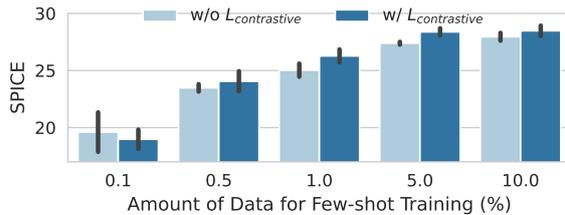


Figure 4: Performance of applying our iNLG on BART-base for few-shot concept-to-text with ablated training objective $\mathcal{L}_{\text{contrastive}}$ on various few-shot settings. Scores are reported with error bars.

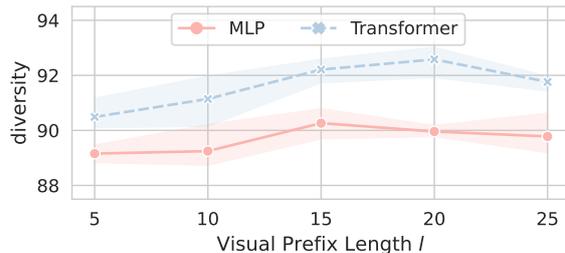


Figure 5: Performance of our iNLG on few-shot sentence completion with various visual prefix lengths and with MLP and Transformer as mapping network. Scores are reported with error bands.

ments, we also tried a simple Multi-Layer Perceptron (MLP) with two fully-connected layers. As shown in Figure 5, the Transformer-based mapping network outperforms MLP on all listed l . MLP has the best performance when visual prefix length $l = 15$, while the Transformer-based mapping network scores highest when $l = 20$.

Model Weight Tuning Table 6 compares the influence of pre-training/tuning the weights of different modules of our iNLG. Generally speaking, tuning the weights during training outperforms freezing the weights, which applies to both the base LM and the mapping network. In addition, considering our few-shot setup, pre-training the mapping network also helps our iNLG gain better performances. The best combination is applying the pre-trained mapping network, and tuning it together with the base LM on the few-shot downstream task.

Showcase Figure 6 provides two showcases on few-shot sentence completion and story generation to compare our iNLG with the GPT2-based baselines. SimCTG and GPT2 tuned with text-only corpus rendering repeated segments, either copying from the input context, or simply repeating themselves. In comparison, our iNLG has minor degeneration and writes coherent sentence endings or stories with more creative details in both tasks.

Tune LM	Pretrain Map.	Tune Map.	diversity \uparrow	MAUVE \uparrow
X	X	X	15.52	0.47
X	X	✓	78.20	33.79
X	✓	X	27.06	1.83
X	✓	✓	76.36	25.15
✓	X	X	87.45	48.06
✓	X	✓	88.68	51.81
✓	✓	X	89.05	55.61
✓	✓	✓	92.68	60.62

Table 6: Performance of our iNLG on few-shot sentence completion with ablated settings on whether to tune the LM, pretrain the mapping network (Pretrain Map.) and tune the mapping network (Tune Map.).

Context: Two girls are standing in a yard wearing cheerleading outfits. A girl...	
Generated Image:	
	GPT2 (text-only) is <u>standing in a yard</u> with a group of girls.
	GPT2 +iNLG is wearing a white bikini and blue shorts.
(a) Sentence Completion	
Context: Barry at the Mall. Barry and his friends spent a day at the mall.	
Generated Image:	
	SimCTG Barry was at the mall with his friends. They were all very excited to see him at the mall. <u>They all had a great time at the mall. They all had a great time at work. They all had a great time at school</u>
	GPT2 (text-only) <u>They went to the local mall and bought a new car. They drove to the mall and bought a new car. They drove to the mall and bought a new car.</u> Barry's friends were very happy they <u>bought a new car.</u>
	GPT2 +iNLG They were swamped by the crowds. Barry was scared to go home. They tried to avoid the mall, but they were stopped by the people at the mall.
(b) Story Generation	

Figure 6: Sentence ending and stories generated by GPT2-based methods tuned with 1% of the training data. *Repetitive contents* are underlined. The sentence ending and story written by our iNLG is coherent with the context, related to the machine-generated image, and has minor degeneration. More demonstrative examples can be found in the Appendix.

6 Conclusion

In this work, we propose iNLG, a framework that introduces machine-generated images to guide open-ended text generation. This endows the machines with the ability of creative visualization that human writers often demonstrate. We distill the vision information from the pre-trained multimodal models and further construct visual prefixes to guide language models to perform text generation with the teacher forcing and the contrastive objective. Extensive experiments show the effectiveness of iNLG in open-ended text generation tasks, including text completion, story generation, and concept-to-text generation in few-shot settings.

Limitations

This work mainly focuses on open-ended text generation, where the search space for the target output is infinite, and the language model would benefit from additional visual imagination distilled from large text-to-image generation models to produce coherent and meaningful content. However, we should note here that despite the commendable performance of text-to-image generation models, there are certain terms and concepts that are inherently challenging to visualize, such as numerical values and abstract philosophical terms. This problem itself is an interesting open research question for all tasks involving text-and-vision.

In our current approach, the images are generated offline. In future work, one may explore the integration of text-to-image and image-to-text modules in an end-to-end manner, which may be more suitable for longer text generation that is not covered in this work.

Text-to-image generation models currently have a length limit on the input text prompt, which may impede their ability to visualize long text inputs in a single image. Furthermore, as previously discussed, text-to-image models may also encounter difficulties in generating images of complex scenes or situations that are challenging to depict through a single image. Future research could explore the use of multiple images or supplementary videos as visual input in order to provide a more comprehensive representation of the scene or situation in question. The iNLG framework can be easily extended to take video representation by taking longer visual prefixes or iteratively applying visual prefixes at each step.

Ethics Statement

In this work, we use pre-trained multimodal models to visualize machine imagination. The machine-generated images may contain uncontrolled bias if any inductive bias exists from the pre-training data. Even though we do not witness such an issue in our study, this may be a potential factor that affects the quality of the generated text. We do not anticipate any major ethical concerns given that all the datasets and models used in this study have already been released to the public. We reproduce baselines with the released code repository. For human evaluation, our study is approved for IRB exempt. The estimated hourly wage paid to MTurk annotators is \$10.

Acknowledgement

The research was sponsored by the U.S. Army Research Office and was accomplished under Contract Number W911NF-19-D-0001 for the Institute for Collaborative Biotechnologies. This work was also supported by the National Science Foundation award #2048122. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

References

- Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2016. [Spice: Semantic propositional image caption evaluation](#). In *ECCV*.
- 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, Ann Arbor, Michigan. Association for Computational Linguistics.
- 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, Hao Tan, and Mohit Bansal. 2021. [Unifying vision-and-language tasks via text generation](#). In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, pages 1931–1942. PMLR.
- Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. [On the properties of neural machine translation: Encoder–decoder approaches](#). In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, pages 103–111, Doha, Qatar. Association for Computational Linguistics.
- Gordon Christie, Ankit Laddha, Aishwarya Agrawal, Stanislaw Antol, Yash Goyal, Kevin Kochersberger, and Dhruv Batra. 2016. [Resolving language and vision ambiguities together: Joint segmentation & prepositional attachment resolution in captioned scenes](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1493–1503, Austin, Texas. Association for Computational Linguistics.
- Katherine Crowson, Stella Rose Biderman, Daniel Kornis, Dashiell Stander, Eric Hallahan, Louis Castricato, and Edward Raff. 2022. [Vqgan-clip: Open domain image generation and editing with natural language guidance](#). In *ECCV*.
- Yann Dauphin, Angela Fan, Michael Auli, and David Grangier. 2017. [Language modeling with gated convolutional networks](#). In *ICML*.
- Desmond Elliott, Stella Frank, Khalil Sima’an, and Lucia Specia. 2016. [Multi30k: Multilingual english-german image descriptions](#). In *Proceedings of the 5th Workshop on Vision and Language, hosted by the 54th Annual Meeting of the Association for Computational Linguistics, VL@ACL 2016, August 12, Berlin, Germany*, pages 70–74. Association for Computational Linguistics (ACL). 5th Workshop on Vision and Language, VL 2016 ; Conference date: 12-08-2016 Through 12-08-2016.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018. [Hierarchical neural story generation](#). In *ACL*, pages 889–898, Melbourne, Australia. Association for Computational Linguistics.
- Linda B Gambrell and Ruby J Bales. 1986. Mental imagery and the comprehension-monitoring performance of fourth- and fifth-grade poor readers. *Reading Research Quarterly*, pages 454–464.
- Linda B Gambrell and Patricia S Koskinen. 2002. [Imagery: A strategy for enhancing comprehension](#). *Comprehension instruction: Research-based best practices*, pages 305–318.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann Dauphin. 2017. [Convolutional sequence to sequence learning](#). In *ICML*.
- Seraphina Goldfarb-Tarrant, Tuhin Chakrabarty, Ralph Weischedel, and Nanyun Peng. 2020. [Content planning for neural story generation with aristotelian rescoring](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4319–4338, Online. Association for Computational Linguistics.
- Michael Grubinger, Paul D. Clough, Henning Müller, and Thomas Deselaers. 2006. [The iapr tc-12 benchmark: A new evaluation resource for visual information systems](#).
- Jian Guan, Fei Huang, Zhihao Zhao, Xiaoyan Zhu, and Minlie Huang. 2020. [A knowledge-enhanced pre-training model for commonsense story generation](#). *ArXiv*, abs/2001.05139.
- Jian Guan, Yansen Wang, and Minlie Huang. 2019. [Story ending generation with incremental encoding and commonsense knowledge](#). In *AAAI*.
- Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Nieves. 2015. [Activitynet: A large-scale video benchmark for human activity understanding](#). *CVPR*, pages 961–970.

- Ting-Hao Kenneth Huang, Francis Ferraro, Nasrin Mostafazadeh, Ishan Misra, Aishwarya Agrawal, Jacob Devlin, Ross Girshick, Xiaodong He, Pushmeet Kohli, Dhruv Batra, C. Lawrence Zitnick, Devi Parikh, Lucy Vanderwende, Michel Galley, and Margaret Mitchell. 2016. [Visual storytelling](#). In *NAACL*, pages 1233–1239, San Diego, California. Association for Computational Linguistics.
- Daphne Ippolito, David Grangier, Douglas Eck, and Chris Callison-Burch. 2020. [Toward better storylines with sentence-level language models](#). In *ACL*, pages 7472–7478, Online. Association for Computational Linguistics.
- Victoria L Joffe, Kate Cain, and Nataša Marić. 2007. [Comprehension problems in children with specific language impairment: does mental imagery training help?](#) *International Journal of Language & Communication Disorders*, 42(6):648–664.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *ACL*, pages 7871–7880, Online. Association for Computational Linguistics.
- Gen Li, Nan Duan, Yuejian Fang, Ming Gong, and Daxin Jiang. 2020. [Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training](#). In *AAAI*, pages 11336–11344. AAAI Press.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. [A diversity-promoting objective function for neural conversation models](#). In *NAACL*, pages 110–119, San Diego, California. Association for Computational Linguistics.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. [Visualbert: A simple and performant baseline for vision and language](#). *ArXiv*, abs/1908.03557.
- Bill Yuchen Lin, Wangchunshu Zhou, Ming Shen, Pei Zhou, Chandra Bhagavatula, Yejin Choi, and Xiang Ren. 2020. [CommonGen: A constrained text generation challenge for generative commonsense reasoning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1823–1840, Online. Association for Computational Linguistics.
- Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. [Microsoft coco: Common objects in context](#). In *ECCV*.
- Ye Liu, Yao Wan, Lifang He, Hao Peng, and Philip S. Yu. 2021. [Kg-bart: Knowledge graph-augmented bart for generative commonsense reasoning](#). In *AAAI*.
- Quanyu Long, Mingxuan Wang, and Lei Li. 2021. [Generative imagination elevates machine translation](#). In *NAACL*, pages 5738–5748, Online. Association for Computational Linguistics.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. [Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks](#). In *NeurIPS*.
- Yujie Lu, Wanrong Zhu, Xin Wang, Miguel Eckstein, and William Yang Wang. 2022. [Imagination-augmented natural language understanding](#). In *NAACL*, pages 4392–4402, Seattle, United States. Association for Computational Linguistics.
- Huaishao Luo, Lei Ji, Botian Shi, Haoyang Huang, Nan Duan, Tianrui Li, Xilin Chen, and Ming Zhou. 2020. [Univlm: A unified video and language pre-training model for multimodal understanding and generation](#). *ArXiv*, abs/2002.06353.
- Kaixin Ma, Filip Ilievski, Jonathan Francis, Satoru Ozaki, Eric Nyberg, and Alessandro Oltramari. 2021. [Exploring strategies for generalizable commonsense reasoning with pre-trained models](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5474–5483, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ron Mokady, Amir Hertz, and Amit H Bermano. 2021. [Clipcap: Clip prefix for image captioning](#). *arXiv preprint arXiv:2111.09734*.
- Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James Allen. 2016. [A corpus and cloze evaluation for deeper understanding of commonsense stories](#). In *NAACL*, pages 839–849, San Diego, California. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *ACL*.
- Krishna Pillutla, Swabha Swayamdipta, Rowan Zellers, John Thickstun, Sean Welleck, Yejin Choi, and Zaïd Harchaoui. 2021. [MAUVE: measuring the gap between neural text and human text using divergence frontiers](#). In *NeurIPS*, pages 4816–4828.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#). In *ICML*.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#).
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.

- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. [Zero-shot text-to-image generation](#). *ArXiv*, abs/2102.12092.
- Robin Rombach, A. Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. [High-resolution image synthesis with latent diffusion models](#). *CVPR*, pages 10674–10685.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. [U-net: Convolutional networks for biomedical image segmentation](#). In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241, Cham. Springer International Publishing.
- Mark Sadoski and Allan Paivio. 2000. [Imagery and text: A dual coding theory of reading and writing](#). Lawrence Erlbaum Associates Publishers.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L. Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, Seyedeh Sara Mahdavi, Raphael Gontijo Lopes, Tim Salimans, Jonathan Ho, David Fleet, and Mohammad Norouzi. 2022. [Photorealistic text-to-image diffusion models with deep language understanding](#). *ArXiv*, abs/2205.11487.
- Jost Schatzmann, Blaise Thomson, Karl Weilhammer, Hui Ye, and Steve Young. 2007. [Agenda-based user simulation for bootstrapping a POMDP dialogue system](#). In *NAACL (Companion Volume, Short Papers)*, pages 149–152, Rochester, New York. Association for Computational Linguistics.
- Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade W Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa R Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. 2022. [LAION-5b: An open large-scale dataset for training next generation image-text models](#). In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Abigail See, Aneesh Pappu, Rohun Saxena, Akhila Yerukola, and Christopher D. Manning. 2019. [Do massively pretrained language models make better storytellers?](#) In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 843–861, Hong Kong, China. Association for Computational Linguistics.
- Haoyue Shi, Jiayuan Mao, Kevin Gimpel, and Karen Livescu. 2019. [Visually grounded neural syntax acquisition](#). In *ACL*, pages 1842–1861, Florence, Italy. Association for Computational Linguistics.
- Anuroop Sriram, Heewoo Jun, Sanjeev Satheesh, and Adam Coates. 2018. [Cold fusion: Training seq2seq models together with language models](#). In *InterSpeech 2018, 19th Annual Conference of the International Speech Communication Association, Hyderabad, India, 2-6 September 2018*, pages 387–391. ISCA.
- Yixuan Su, Tian Lan, Yahui Liu, Fangyu Liu, Dani Yogatama, Yan Wang, Lingpeng Kong, and Nigel Collier. 2022a. [Language models can see: Plugging visual controls in text generation](#). *ArXiv*, abs/2205.02655.
- Yixuan Su, Tian Lan, Yan Wang, Dani Yogatama, Lingpeng Kong, and Nigel Collier. 2022b. [A contrastive framework for neural text generation](#). In *NeurIPS*.
- Chen Sun, Austin Myers, Carl Vondrick, Kevin P. Murphy, and Cordelia Schmid. 2019. [Videobert: A joint model for video and language representation learning](#). *ICCV*, pages 7463–7472.
- Ben Swanson, Kory Mathewson, Ben Pietrzak, Sherol Chen, and Monica Dinalescu. 2021. [Story centaur: Large language model few shot learning as a creative writing tool](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 244–256, Online. Association for Computational Linguistics.
- Hao Tan and Mohit Bansal. 2020. [Vokenization: Improving language understanding with contextualized, visual-grounded supervision](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2066–2080, Online. Association for Computational Linguistics.
- Aäron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. [Representation learning with contrastive predictive coding](#). *ArXiv*, abs/1807.03748.
- Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2015. [Cider: Consensus-based image description evaluation](#). *CVPR*, pages 4566–4575.
- PeiFeng Wang, Jonathan Zamora, Junfeng Liu, Filip Ilievski, Muhao Chen, and Xiang Ren. 2022a. [Contextualized scene imagination for generative commonsense reasoning](#). In *ICLR*.
- Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. 2022b. [Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework](#). *ICML*.
- Zhongyu Wei, Qianlong Liu, Baolin Peng, Huaixiao Tou, Ting Chen, Xuanjing Huang, Kam-fai Wong, and Xiangying Dai. 2018. [Task-oriented dialogue system for automatic diagnosis](#). In *ACL (Volume 2: Short Papers)*, pages 201–207, Melbourne, Australia. Association for Computational Linguistics.
- Sean Welleck, Ilia Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2020. [Neural text generation with unlikelihood training](#). In

ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020. OpenReview.net.

Tsung-Hsien Wen, Milica Gašić, Nikola Mrkšić, Pei-Hao Su, David Vandyke, and Steve Young. 2015. [Semantically conditioned LSTM-based natural language generation for spoken dialogue systems](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1711–1721, Lisbon, Portugal. Association for Computational Linguistics.

Tsung-Hsien Wen, David Vandyke, Nikola Mrkšić, Milica Gašić, Lina M. Rojas-Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2017. [A network-based end-to-end trainable task-oriented dialogue system](#). In *ACL*, pages 438–449, Valencia, Spain. Association for Computational Linguistics.

Jie Wu, Ian Harris, and Hongzhi Zhao. 2021. [Spoken language understanding for task-oriented dialogue systems with augmented memory networks](#). In *NAACL*, pages 797–806, Online. Association for Computational Linguistics.

Ning Xie, Farley Lai, Derek Doran, and Asim Kadav. 2019. [Visual entailment: A novel task for fine-grained image understanding](#). *ArXiv*, abs/1901.06706.

An Yan, Zexue He, Xing Lu, Jiang Du, Eric Chang, Amilcare Gentili, Julian McAuley, and Chun-Nan Hsu. 2021. [Weakly supervised contrastive learning for chest X-ray report generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4009–4015, Punta Cana, Dominican Republic. Association for Computational Linguistics.

Pengcheng Yang, Boxing Chen, Pei Zhang, and Xu Sun. 2020. [Visual agreement regularized training for multi-modal machine translation](#). In *AAAI*, volume 34, pages 9418–9425.

Lili Yao, Nanyun Peng, Ralph M. Weischedel, Kevin Knight, Dongyan Zhao, and Rui Yan. 2019. [Plan-and-write: Towards better automatic storytelling](#). In *AAAI*.

Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [HellaSwag: Can a machine really finish your sentence?](#) In *ACL*, pages 4791–4800, Florence, Italy. Association for Computational Linguistics.

Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *ICLR*.

Wangchunshu Zhou, Dong-Ho Lee, Ravi Kiran Selvam, Seyeon Lee, and Xiang Ren 0001. 2021. [Pre-training text-to-text transformers for concept-centric common sense](#). In *ICLR 2021*. OpenReview.net.

Wanrong Zhu, Xin Eric Wang, An Yan, Miguel P. Eckstein, and William Yang Wang. 2021. [Imagine: An imagination-based automatic evaluation metric for natural language generation](#). *ArXiv*, abs/2106.05970.

A Appendix

A.1 Experiment Details

Pretraining We pre-train the mapping network for GPT-2-base (Radford et al., 2019) on the MSCOCO (Lin et al., 2014) dataset with 414,113 (image, text) pairs for training. We pre-train the mapping network for BART-base (Lewis et al., 2020) on VIST (Huang et al., 2016) story-in-sequence subset, with 141,593 (image, text) pairs for training after excluding the images that the users have removed. For each setting, we pre-train the mapping network for 5 epochs with a batch size of 128, learning rate of 2e-5, weight decay of 0.01, and warmup steps of 5,000.

Few-Shot Training for Downstream Tasks Table 7 lists out the hyperparameters we used during few-shot experiments on the three open-ended text generation tasks.

Hyperparameters	Concept-to-Text	Text Completion	Story Generation
Base LM	BART-base	GPT2-base	GPT2-base
Batch Size	8	8	8
Training Epoch	20	20	20
$N_{\text{no_contra}}$	4	10	15
λ	1.5	1	0.2
Learning Rate	2e-5	2e-5	2e-5
Weight Decay	0.01	0.01	0.01
Warmup Steps	400	400	400
Max Output Length	64	100	150
Num of Beam	10	10	10

Table 7: Hyperparameter settings for few-shot open-ended text generation.

Parameter Search We tried the learning rate in the following setting: {1e-5, 2e-5, 5e-5, 1e-4}, and tried the batch size in {4, 8, 16, 32}.

Parameter Size Table 8 lists out the parameter size for the network modules used in our study.

Environment & Run Time Table 9 lists out the execution time for the three open-ended text generation tasks with 1% of the training data. Experiments are conducted on NVIDIA A100.

A.2 Human Evaluation

We invite Amazon Mechanical Turk¹² annotators to judge the quality of the generated text. Figure 7 shows an example template we use for head-to-head comparison.

¹²<https://www.mturk.com/>

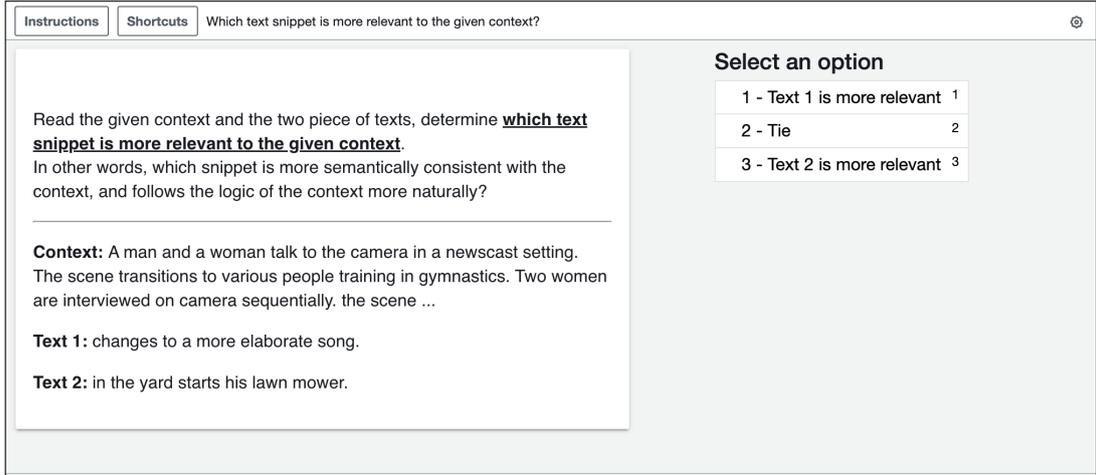


Figure 7: A screenshot of the Amazon Mechanical Turk interface for our human evaluation on text coherency.

Task	Model	Parameter Size
Sentence Completion	StoryEndGen	11M
	GPT-2 base	117M
	GPT-2 base+iNLG	160M
Story Generation	Action-Plan	43M
	Plan-and-Write	34M
	SimCTG	117M
Concept-to-Text	BART-base	110M
	KGBART	439M
	ModelAdapt	110M
	Imagine-and-Verbalize	880M
	BART-base+iNLG	153M

Table 8: Parameter sizes of the network modules used in our study.

Dataset	Text-only	+ iNLG
ActivityNet	50min	70min
ROCStories	70min	95min
CommonGen	40min	55min

Table 9: The average execution time for one single run (training + inference) on each dataset. Text generation experiments are conducted on NVIDIA A100.

A.3 More Showcases

Figure 8 compares the images retrieved from Yahoo Image Search and the images generated by StableDiffusion-v1-1 (Rombach et al., 2022), which is the text-to-image generation model we used in this work. Figure 9 and Figure 10 show more examples comparing the sentence endings and stories generated by different models.

Context 1: One of the guys hits the ball over to the other side and they hit it back. Then on the other side of the beach there is a group of women also playing volleyball. They...

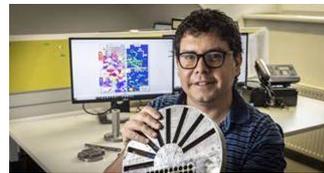


(a1) Retrieved Image



(b1) Generated Image

Context 2: A boy is talking to a camera. He goes into a bathroom and drinks a cup of mouthwash. He...



(a2) Retrieved Image



(b2) Generated Image

Figure 8: With the context as input, (a1)(a2) is the first returned image by the Yahoo image search engine,¹³ and (b1)(b2) is generated by StableDiffusion-v1-1 (Rombach et al., 2022). The two input contexts are from the ActivityNet (Heilbron et al., 2015) subset in Hel-laSwag (Zellers et al., 2019).

¹³The screenshots of the search results returned by Yahoo Image Search as of Feb.3rd 2023: [link1](#), [link2](#).

Context: A leaf blower is shown blowing a large pile of leaves across a green lawn in front of residential houses. The leaves...

Generated Image:



StoryEndGen	is her hands . the woman
GPT2 (no finetune)	are then blown by a small wind turbine.
GPT2 (text-only finetune)	are then shown in a large circle and the roof is shown in a close up.
GPT2 + iNLG	are placed on the ground and the man is shown sitting on the back.

(a)

Context: Men are standing in the edge of a trampoline preparing to make a jump into a pool in a roofed pool. People...

Generated Image:



StoryEndGen	then then the camera and then the camera and the camera and the camera
GPT2 (no finetune)	are standing in the edge of a trampoline preparing to amke a jump into a pool in a roofed pool.
GPT2 (text-only finetune)	are standing in the middle of the pool preparing to jump into a pool in a roofed pool.
GPT2 + iNLG	are swimming in the pool and throwing a ball.

(b)

Context: They mix the eggs around a bowl and place butter and milk into another bowl and mix them all together. They...

Generated Image:



StoryEndGen	the cake ups and the cake and then the cake and then the cake and the cake and then the cake and then the cake and
GPT2 (no finetune)	will be very soft and fluffy.
GPT2 (text-only finetune)	are ready to use.
GPT2 + iNLG	then put a bowl on the end of the sink and put the mixture in the sink.

(c)

Figure 9: Comparisons on few-shot sentence completion performance on ActivityNet.

Context: Sushi. Jerry loves to eat food from other nations.

Generated Image:



Action-Plan	She was to to to a . He . He was to to to the . He . She was to to to . He . She was to the the .
Plan-and-Write	the man was a lot of her . good he had to go good his new phone good friends # wanted # made # found # happy good good # home good good good # time good good good # bought # time
SimCTG	Tom was hungry. He went to the sushi place. There was a lot of sushi. Tom was hungry. He went to the sushi place
GPT2 (no finetune)	He's a big fan of the Japanese.
GPT2 (text-only finetune)	He goes to the country to buy the freshest seafood. He buys a lot of seafood. He buys more seafood and buys more seafood. Jerry buys more seafood and eats better.
GPT2 + iNLG	He was unable to find a slice he wanted. He grabbed a slice of sushi at a fast food restaurant. He was able to eat it and it was the perfect meal.

(a)

Context: Hole in the House. Tim was putting away some paperwork in the attic.

Generated Image:



Action-Plan	He was to to to the . He . He was to to the . . He was to to to the . . He was to the the .
Plan-and-Write	# i was a very hot . </s> he decided good to the store good <EOT> friends # wanted # made # found # happy <EOL> # my friend went good with good , she had good and looked in good # time # found # home
SimCTG	My neighbor's house was burglarized. I went to investigate. My neighbor's house was burglarized. My neighbor's house was burglarized. My neighbor's house was burglarized
GPT2 (no finetune)	He was trying to figure out how to get the house to be ready for the next day. He was trying to figure out how to get the house to be ready for the next day. "I'm not going to be able to do that," he said. "I'm not going to be able to do that."
GPT2 (text-only finetune)	He was trying to write a letter to his boss. He was trying to get his boss to write a letter to him. Tim was frustrated.
GPT2 + iNLG	He saw a map of the area. He went to the bathroom to check. There was nothing there. He was surprised to see it was a loophole.

(b)

Figure 10: Comparisons on few-shot story generation performance on ROCStories.