# **Image Content Generation with Causal Reasoning**

Xiaochuan Li<sup>1,4</sup>, Baoyu Fan<sup>2,1,\*</sup>, Runze Zhang<sup>1</sup>, Liang Jin<sup>1</sup>, Di Wang<sup>1</sup> Zhenhua Guo<sup>1</sup>, Yaqian Zhao<sup>1</sup>, Rengang Li<sup>3,1</sup>

<sup>1</sup>Inspur Electronic Information Industry Co.,Ltd.

<sup>2</sup>Nankai University

<sup>3</sup>Tsinghua University

<sup>4</sup>Shandong Massive Information Technology Research Institute

lixiaochuan2088@gmail.com, fanbaoyu@foxmail.com, {zhangrunze, jinliang, wangdi11, guozhenhua}@ieisystem.com zhaoyaqian@ieee.org, lirengang.hsslab@gmail.com

#### Abstract

The emergence of ChatGPT has once again sparked research in generative artificial intelligence (GAI). While people have been amazed by the generated results, they have also noticed the reasoning potential reflected in the generated textual content. However, this current ability for causal reasoning is primarily limited to the domain of language generation, such as in models like GPT-3. In visual modality, there is currently no equivalent research. Considering causal reasoning in visual content generation is significant. This is because visual information contains infinite granularity. Particularly, images can provide more intuitive and specific demonstrations for certain reasoning tasks, especially when compared to coarse-grained text. Hence, we propose a new image generation task called visual question answering with image (VQAI) and establish a dataset of the same name based on the classic Tom and Jerry animated series. Additionally, we develop a new paradigm for image generation to tackle the challenges of this task. Finally, we perform extensive experiments and analyses, including visualizations of the generated content and discussions on the potentials and limitations. The code and data are publicly available under the license of CC BY-NC-SA 4.0 for academic and non-commercial usage at: https://github.com/IEIT-AGI/ MIX-Shannon/blob/main/projects/VQAI/lgd\_vqai.md.

#### Introduction

AI-generated content (AIGC), also known as generative AI (GAI), recently gained a surge of development (Zhang et al. 2023a,b; Cao et al. 2023; Balaji et al. 2022), covering several areas such as image (Ramesh et al. 2021, 2022; Saharia et al. 2022; Yu et al. 2022; Rombach et al. 2022), text (Raffel et al. 2020; Radford et al. 2018, 2019; Brown et al. 2020; OpenAI 2023; Vinyals et al. 2015), 3D (Fu et al. 2022; Jahan, Guan, and Van Kaick 2021; Liu et al. 2022; Mildenhall et al. 2021), and speech (Qian et al. 2014; Ze, Senior, and Schuster 2013; Zen and Sak 2015). Since ChatGPT emerged, people have been amazed by its performance while recognizing the reasoning potential in the generated text content (Bang et al. 2023). In particular, some recent studies have started to delve into the text's reasoning ability (Kojima et al. 2022;

Wang et al. 2023; Wei et al. 2022; Shum, Diao, and Zhang 2023), including causal reasoning (Wang et al. 2022; IDEA-CCNL 2021), in GAI.



Figure 1: Generated results of Stable Diffusion v2.1. The generated results strictly follow the guidance of the text, ignoring other content caused by the implied conditions.

However, majority of these works have primarily focused on text content generation, with only a limited number of studies exploring other modalities like images. Although some studies have tried to use images as input and achieved good output results, except for some scalable vector graphics (SVG) representation for sketches or doodles (OpenAI 2023), this field has been scarcely studied. AIGC is currently evolving towards making the generated content more realistic. More specifically, these generated contents cover as many requirements as possible in the guidance and present more realistic details that amaze the human eye or ear. However, these generative models are difficult to follow when underlying cause-and-effect logic is implicit in prompts like an implied condition or relationship between objects.

Regarding the image AIGC solely, popular models do not exhibit satisfying reasoning abilities. As shown in Figure 1,

<sup>\*</sup>Corresponding author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



Figure 2: Task definition of VQAI.

the model is competent when we ask for a broken glass. Once the "filled with orange juice" condition is added, a hidden fact is that "the juice will be spilled out" since "the glass breaks". Unfortunately, the generation fails. However, it can be generated smoothly if we write this fact obviously into the guidance, as shown in the lower left corner. Finally, we give only the events without any prompts of possible outcomes. The generated image does not even include the "broken" statement.

In the field of visual content generation, it is valuable to consider the ability of causal reasoning during image generation. Specifically, since images contain information at infinite granularity, images can give a more intuitive and specific demonstration for some inference tasks, especially in comparison with the coarse-grained text. In this study, we consider image content generation with causal reasoning. Thus, we propose a new task as shown in Figure 2.

In particular, this task is somewhat similar to the imageediting task in terms of the form of input and output. More specifically, the difference is that we do not provide an exact description of the differences used to edit, but rather a text question containing a condition. Furthermore, models need to generate appropriate image content based on the implicit cues in the text and initial image. From this perspective, our work can also be seen as an extension of the classical multimodal task visual question answering (VQA) (Antol et al. 2015) in terms of output modality. Thus we also refer to this task as visual question answering with image (VQAI).

Accordingly, we make a new dataset for this task based on the classic *Tom and Jerry* cartoon series for two main reasons. First, *Tom and Jerry* has a more straightforward worldview than the natural scenarios, meaning that the causal relationship between characters is more straightforward and clear, with very few indirect causal events like *emotional hiding* and *complex strategies*. Second, it gives more prominence to the visual aspect of behaviors. Particularly, it weakens speech as much as possible and describes relationships through visual states like movements and expressions. Besides, due to the animation, the variations of objects and backgrounds are relatively controllable, facilitating our first attempt at this task.

Particularly, we develop a new method for this task. An obvious idea is concatenating a multimodal comprehension module and a visual generator. While the former is used to generate sensible text, the latter performs image editing based on the former's output. The text acts as a bridge in this pipeline. However, this exposes a considerable risk - there is so much less information in text than in an image.

Consequently, it would take an enormous amount of text to replace the content in the image. This is most likely beyond the comprehension capability of the editing model and even beyond the token length limit. Moreover, making the language model generate long enough text is complex. Therefore, we propose a hidden space-guided causal image generation method and conduct extensive experiments on the proposed dataset to demonstrate the scheme's effectiveness.

We summarize our main contributions as follows:

- We rethink image AIGC with causal reasoning and propose a new task called VQAI.
- Additionally, a new dataset is proposed to support the study of causal image generation.
- Furthermore, we analyze the challenges of this task and propose a new approach to solve it. Extensive experiments demonstrate the effectiveness of our method.

## **Related Work**

Image AIGC: Image generation is an important research area of visual AIGC that drawn huge interest among researchers (Zhang et al. 2023a,b; Cao et al. 2023; Balaji et al. 2022). In recent years, underlying generative models have been continuously proposed to promote development in this field. Variational auto-encoder (VAE) (Kingma and Welling 2013) is an auto-encoder that learns data distribution from latent space, and it can change the generated image by verifying the input encoding. The generative adversarial network (GAN) (Goodfellow et al. 2020; Dhariwal and Nichol 2021) trains a discriminator and a generator, respectively, based on the deep network to achieve automatic image generation, driving a wave of research trends. PixelRNN (Van Den Oord, Kalchbrenner, and Kavukcuoglu 2016) generates reasonable visual patches based on prior pixel sequence. More recently, the diffusion model (Ho, Jain, and Abbeel 2020) learns information degradation due to noise and generates images systematically using the learned patterns.

Meanwhile, in the past two years, text-guided image generation (text-to-image) has become popular with the rise of multimodal research. DALL-E (Ramesh et al. 2021) uses a pre-trained discrete variational auto-encoder (dVAE) to extract tokens for the image and an auto-regressive transformer to generate the image. Stable/Latent Diffusion (Rombach et al. 2022) replaces the image with encoded features as a supervised signal and restores them to images via a visual decoder. Furthermore, eDiffi (Balaji et al. 2022) generates better images by integrating expert denoisers in the diffusion model, using both CLIP (Radford et al. 2021) and T5 (Raffel et al. 2020) as text encoders.

Meanwhile, some other tasks for image content generation are also derived. Image editing aims to edit an image according to a given text or another description form. Imagic (Kawar et al. 2022) performs various text-based semantic editing on a single image, including highly complex nonrigid changes like pose changes and editing multiple objects. InstructPix2Pix (Brooks, Holynski, and Efros 2022) introduces small structural changes to the diffusion model and fine-tunes it to gain editing capabilities. More so, the story continuation task generates subsequent images based on the initial image and plot synopsis. StoryDALL-E (Maharana, Hannan, and Bansal 2022) uses a pre-trained text-to-image transformer to make the plot of the generated content more coherent. AR-LDM (Pan et al. 2022) trains a hidden diffusion model to improve the generated images' story continuity and content consistency.

These works promote research in visual AIGC, where image content becomes increasingly controllable and realistic. However, these tasks require images to be generated strictly following text guidance, ignoring the ability to reason during generation. Therefore, this work aims to develop such a topic and investigate causal image content generation.

Reasoning in GAI: Recently, with the birth of ChatGPT, the research for large language models (LLMs) continued to grow in popularity. Particularly, many researchers have started to explore the reasoning abilities embedded in LLMs. It is found that adding encouragement to prompts drives reasoning in the generated text (Wei et al. 2022). Zero-Shot-CoT (Kojima et al. 2022) concatenates "Let's think step by step" after a question to get a more detailed reasoning step and achieves better performance on question answering (QA) tasks. Manual-CoT (Wei et al. 2022) manually designs a few question-answer samples to guide the language model to continue the chain of thoughts. Automatic-CoT (Shum, Diao, and Zhang 2023) constructs a candidate pool of rationale chains based on a small-labeled dataset and selects the best combination for CoT prompting by employing a variance-reduced policy gradient strategy. Besides, some researchers have started to analyze the causal/counterfactual reasoning ability embodied in LLMs, such as Randeng-Deduction (Wang et al. 2022; IDEA-CCNL 2021). These works illustrate that LLM-based GAI exhibits interesting reasoning capabilities, at least in text generation.

In the multimodal domain, MM-CoT (Zhang et al. 2023c) transfers the CoT to image-text samples and enables detailed rationale and answer generation by fine-tuning an LM. GPT4 (OpenAI 2023) demonstrates the results of causal reasoning with images and can even generate scribbles of simple images in SVG representation. However, although these works considered images in GAI, they did not include images as outputs. This study refers to related works in the field of LLM and multimodality to further investigate the causal capabilities in image generation.

## **Visual Question Answering with Image**

This section presents the dataset for the proposed task VQAI. We make the code to access the dataset publicly available under a CC BY-SA 4.0 license. More details are released in *Supplementary Material* (Li et al. 2023b).

#### **Task Definition**

*Syllogism* (Smiley 1973) is a basic unit of causal reasoning and is divided into three parts: *major premise, minor premise* and *conclusion*. The *major premise* is the statement of a general or universal nature. The *minor premise* is the statement about a particular case. The *conclusion* is a corollary to accepting the premises. To study causal reasoning in the visual task, we use this *syllogism* form to formulate the task. In particular, as shown in Figure 3 (a), we construct a sample comprising three parts: i) an initial image as the major premise, which is used to describe the relationship between the current scene and the objects; ii) an interrogative/question as the minor premise, containing a causal condition in the current scenario; and iii) an answer image as the conclusion describing a reasonable result considering both constraints. This formulation is like a sample of VQA; thus, our work can also be seen as an extension of VQA considering causal reasoning on the output modality.

### **Data Collection**



Figure 3: A sample from the VQAI dataset.

For the prompt VQAI task, we produce a new dataset. All pictures used in the dataset are sampled from the Tom and Jerry cartoon series. We adopt this cartoon for several main reasons. First, compared to the complex scenes of nature, the worldview of cartoons is often simplified. Specifically, the relationships between different entities are greatly simplified in the world of Tom and Jerry, and there are very few overly complex events such as "emotional hiding" and "complex strategies". Instead, the exaggerated drawing style usually highlights or emphasizes the character's reaction to a particular condition. This means that the cause-and-effect relationship between characters is more straightforward and clear, facilitating our analysis and study of the task. Second, Tom and Jerry attaches great importance to the presentation of visuals. Compared with some other animated films, this cartoon ignores the necessity of language as much as possible. Particularly, it tends to convey the moods and reactions of its characters through expressions, movements, or states. This is beneficial to our exploration of image causal generation. Besides, as an animated film, the changes in backgrounds, characters, and objects are relatively controllable. This facilitates our first attempt at this novel task. Meanwhile, it reduces the difficulty of data collection.

Specifically, we download 755 episodes of *Tom and Jerry* from public sources, hand-crop pairs of images where causal relationships exist, and label the interrogative sentences that contain the conditions.

### Annotations

*Causal Questions:* The annotators are asked to annotate a causal question of the text on each pair of causal images.



Figure 4: Demonstration and proportion of the five categories of samples in the dataset.

Particularly, the question needs to give conditions under which an initial image evolves rationally, and the answer image must be a reasonable result of the evolution of the initial image under this condition. An example is shown in Figure 3 (a). In detail, each image pair is first checked by the annotator for the existence of valid causal relationships, and then these relationships are summarized in a single question, such as "What happens if [condition]?", "What will it do if [condition]?" and so on.

*Causal Chain:* The annotation of causal questions completes the *syllogism.* However, this task is still challenging, especially as some events develop, requiring multiple reasoning steps. It is indeed difficult for image generation models to answer complex causal questions (which is analyzed further in Section ). Moreover, models cannot be assessed whether they truly learn causal reasoning capabilities or just fit the statistical bias on the dataset. Therefore, we annotate the steps of reasoning, called causal chains, for a part of the samples. It shows how the first image develops systematically into the second under the conditions given in the causal question. In the causal chain, each edge represents one inference step, and each node represents a variation of the event development, as shown in Figure 3 (b).

In allow the inference process to be better structured, we classify edges and nodes separately. Specifically, edges are classified into two types: i) to express causal reasoning, for example, "Jerry kicks dog's foot" **causes** "the dog feels pain", which is conventional forward reasoning; ii) to express the condition or need, for example, "Jerry kicks dog's foot" **needs** "Jerry jumps up", which looks more like a reverse thinking process out the necessary conditions. In addition, nodes are also divided into two types, which are i) **visible** and ii) **invisible** in the image. For example, "feeling pain" is actually a mental activity that is not visible. However, it leads to a "painful expression", which is visible.

## **Quality Control**

We follow strict control rules to ensure the quality of the dataset, reflected in two main aspects – *annotation guidance* and *annotation checking*. In particular, this quality control procedure applies to the causal questions since this part of the data is labeled by an external worker due to its large vol-

ume. Conversely, the causal chains are all labeled by three experienced researchers.

Annotation Guidance: We provide annotators with five strict templates to select image pairs from the video and write causal questions. In short, given an image pair, it is valid when and only if it satisfies one of the following rules:

- Scenery Variation: the scene or environment is modified, such as changes in weather, brightness, and season.
- More Entities: the scene has not been modified, but one or more entities have been added.
- Fewer Entities: the scene has not been modified, but one or more entities have been reduced.
- Entities Variation: the modifications to the scenario are minor with no additions or subtractions of entities.
- Emotion Variation: one or more characters' emotions change, accompanied by expressions or movements.

Figure 4 (a) shows specific examples of these five sample categories, whose proportions are represented as shown in Figure 4 (b).

Annotation Checking: The researchers review each causal question label upon submission consistent with the above criteria, and samples with unreasonable causal questions are rejected. Ultimately, VQAI contains 17,524 samples, of which 3,809 sets include causal chain annotations.

## Latent Guided Diffusion Model via Frozen Large Language Model

#### Latent Guided Image Generation

We think about causal reasoning and image content generation. Specifically, one of the most straightforward solutions is to use an off-the-shelf image editing model, such as InstructPix2Pix, if only the matching of the model to the data structure is considered. We refer to this approach as question-guided generation, as shown in Figure 5 (a). However, since the input text does not contain explicit information for modification, it is risky to use only an image editor. It may not be able to causal reason. Therefore, we consider cascading a reasoning module before the generator.

The reasoning ability of large language models (LLMs) is widely recognized. In the multimodal domain, some approaches have inserted adapters in LLM and verified that



Figure 5: Three paradigms for causal image generation.

LLM maintains reasoning ability over multimodality on tasks such as VQA and image captioning (Li et al. 2023a). Therefore, it is worth borrowing the inference capability of LLM in the causal image generation process. In this work, we consider two different paradigms, as shown in Figure 5 (b) and (c), and refer to them as answer-guided generation and latent-guided generation, respectively. One is to use LLM to reason about the textual output for multimodal inputs and use that answer to guide the editing model. However, since images are far more information-rich than text, this may introduce new risks. It requires considerable textual description to replace the equivalent amount of information in an image. This may exceed the editing model's comprehension capability and even the token length limit. In this study, we propose a new generative paradigm that uses the encoding features of LLM to guide the generation model.

We propose a new method based on LLM with a diffusion model called latent guided diffusion (LGD), and the method structure is shown in Figure 6. For the use of LLM, we introduce the Q-Former in the BLIP2 (Li et al. 2023a). Q-Former initializes a set of fixed-length query tokens and translates the image information into features that the LLM encoder can read by making cross attention to the image features. These features are concatenated together with the embeddings of language prompts (interrogatives or other forms of instances) to implement different downstream tasks. As shown in the figure, we use the same form of extracting multimodal features for images and causal interrogatives and use them to guide image decoding. On the image decoding side, we refer to the related work of Stable Diffusion (Rombach et al. 2022) and InstructPix2Pix (Brooks, Holynski, and Efros 2022) to fuse the features used for guidance into different stages of UNet through the attention mechanism.

However, this brings new challenges. First, the diffusion model does not recognize the output features of LLM. Since it is costly to construct a large causal dataset, we intend to refrain from training either of them from sketch. Therefore, a space translation for latent features is necessary. It adapts the feature's dimensionality and semantics of the feature. In the structure, we add a fully connected layer in front of the latent diffusion model to translate the distribution of its input.

Moreover, for causal content generation, the output image is essentially a prediction of the subsequent of the initial image under a particular condition. In other words, the features used for guidance need to express the information after being predicted. Unfortunately, we cannot be certain that the output of the LLM encoder can explicitly contain this. Usually, these features are inputted into the LLM decoder, and the prediction is made gradually by autoregression. Therefore, we add a new module for predictive encoding to predict the subsequent steps triggered by causal interrogatives. We realize this through the inspiration of predictive coding (PC) (Aitchison and Lengyel 2017; Huang and Rao 2011; Oord, Li, and Vinyals 2018). In sequence-prediction tasks such as speech, ordered image, and video, PC predicts the prediction space after a given moment based on a given sequence by adding a new series of fully connected layer combinations. This is very similar to the task of causal image prediction, so we transform the output of the LLM encoder into predicted information by setting up fully connected layers and encoding the prediction.

#### **Contrast Causal Predictive Coding**

Contrast predictive coding (CPC) (Oord, Li, and Vinyals 2018) takes the speech or other ordered fragments as input to the encoder to extract ordered features. The features are input to multiple isomorphic predictive coding networks with different weights to obtain predictive features for multiple moments after that fragment. This approach induces potential space to capture valuable information to predict future samples. Ultimately, CPC optimizes the model parameters by constraining the predictive features to the ground truth of the corresponding batch's corresponding moments.

We refer to this form because causal image generation can also be seen as a prediction task. Moreover, the labeled form of the causal chain appears to satisfy the conditions to construct the loss function. However, two risks arise. First, the time-series samples to which this predictive encoding applies are uniform. In other words, the distance between any adjacent frames in the sequence is the same as in a speech or video sequence with a fixed time interval. However, such uniformity does not exist in causal inference. In particular, it is not guaranteed that all neighboring nodes of a causal chain express an equal number of inference steps between them. As in Figure 3 (b), we consider that "mouse kicks dog's foot" (node B) causes "dog feels pain" (node C), followed by "dog shows a painful expression" (nodeD). However, it seems reasonable to derive D directly from B. People express the inference differently, so the causal chain is not uniform. Moreover, the annotation of causal chains has a significant long-tail effect. This may lead to insufficient training of the latter fully-connected layers in traditional PC.

Therefore, we propose contrast causal predictive coding (CCPC). First, we replace several FCs in CPC with one, which is only used to encode whether there is a causal relationship between the two. Specifically, while calculating the loss, we take positive samples from the causal chain of the current sample and select several nodes from other samples as negatives, and optimize the model parameters using contrast learning.



Figure 6: Architecture of Latent Guided Diffusion Model via Frozen Large Language Model.



Figure 7: Structures of CPC (a), CCPC (b), and MCCS (c).

## **Multimodal Causal Chain Supervision**

To enable better characterization of LLM-encoded features, introduce the supervision of the causal chain text generation. Inspired by a work related to chain-of-thoughts (Kojima et al. 2022; Wei et al. 2022; Zhang et al. 2022; Shum, Diao, and Zhang 2023; Zhang et al. 2023c), we use text in the form of causal chains to provide supervised signals. In the training phase, we generate text labels for samples that include causal chain annotations. In particular, a fixed template guidance helps generate this part of the labels as shown in Figure 7(c). Specifically, these text labels provide an auxiliary optimization for the trainable parameter part of the model, as shown in Figure 7 (c).

## **Experiments**

In this section, we show and analyze the experimental phenomena. First, we compare the results of the three generative paradigms in Figure 5. After that, we block latent space translation (LST), CCPC, and multimodal causal chain supervision (MMCS) and analyze the results of the ablation experiments. All experiments are run on an A100 $\times$ 8 server. In the dataset, we divided 17,524 samples into 15,524, 1,000 and 1,000, corresponding to the training, validation and testing sets. Among them, 3809 samples in the training set include causal chain annotations. Regarding the model, the LLM in this work references T5-XXL, and the image decoder uses stable diffusion. In the training phase, we use Flan-T5-XXL (Raffel et al. 2020; Chung et al. 2022) with the original stable diffusion to initialize the parameters. All initial learning rates are set to 3e-5. In the comparison experiments, we use ADAM (Kingma and Ba 2014) as the optimizer. We set the batch size to 16 and the epoch to 20.

Additional, we show more details and analysis in Supple-

mentary Material(Li et al. 2023b).

#### **Evaluation Metrics**

We design CLIP-based (Gal et al. 2022; Radford et al. 2021) and human-based evaluation matrics, respectively. Specifically, we compute the similarity of CLIP features between the generated image and the ground truth, denoted as  $Sim_{Avg}$ . However, given the diversity that results from causal reasoning, it is not reasonable to conclude that a result different from GT is wrong. Thus, we propose  $Sim_{Best@k}$ to compute the maximum value of similarity among the kresults generated. In our experiments, we set k to 9. After that, we introduce AUC based on CLIP score to observe the semantic accuracy of the generated pictures, denoted as  $AUC_{Avg}$  and  $AUC_{Best@k}$ . In addition, we incorporate human evaluations in order to accommodate the diversity of results. We invite 10 researchers to evaluate whether the generated images are semantically causally related to the input to obtain the accuracy. In addition, we ask each evaluator to subjectively select the one they think is the best to compare the generative performance of the different methods, which is denoted as ChosenRate.

## **Causal Image Generation**

We evaluate the three paradigms represented in Figure 5, as shown in Figure 8. It can be clearly observed that the results of the question-guided diffusion model (QGD) are confusing. In particular, the image decoder may incorrectly add something from the question into the image rather than understanding the result to which this question would lead. As shown in the first row of Figure 8, QGD incorrectly generates the content of the word "biscuit" instead of the rat's reaction to losing it. This may be because the decoder does not have the ability to reason. It can understand certain elements or variations that appear in the text and present them in the image modification process while ignoring those implicit inferences from the text. More so, it makes this paradigm easily adaptable to tasks like image editing rather than causal content generation. However, AGD is an improvement of QGD. Since QGD lacks the ability of causal reasoning, a possible solution is to cascade a text reasoning model before the image decoder, whose duties are simultaneously reduced to a single editing model. Fortunately, the examples in Figure 8 show that AGD is effective. Furthermore, The Latentguided diffusion model (LGD) is a further improvement on The Thirty-Eighth AAAI Conference on Artificial Intelligence (AAAI-24)



#### Figure 8: Visualization of the results generated by QGD, AGD, and LGD.



Question: If there is delicious food in front of them, what will be left in the scene in the next moment?

Figure 9: Visualization of the ablation results.

the AGD. This suggests that language does overlook some imperceptible variations in images that may be preserved in the hidden space, as shown in the second row of Figure 8. In addition, we evaluate the three methods quantitatively, as shown in Table 1. The results of all the experiments can be seen in Table 1, where LGD is superior in all metrics.

Methods	QGD	AGD	LGD
$Sim_{Avg}$ (CLIP)	0.8361	0.8444	0.8589
$Sim_{Best@9}$ (CLIP)	0.8831	0.8867	0.9038
$AUC_{Avg}$ (CLIP)	0.8311	0.8394	0.8539
AUC <sub>Best@9</sub> (CLIP)	0.8781	0.8819	0.8987
Acc (human)	0.1695	0.1852	0.3239
ChosenRate (human)	0.1601	0.2310	0.5135

Table 1: Quantitative comparison of three paradigms. CLIP and human evaluations of the mentioned methods.

#### **Ablation Study**

We conduct experiments to analyze the effects of three modules proposed before: latent space translation (LST), contrast causal predictive coding (CCPC), and multimodal causal chain supervision (MCCS). Figure 9 presents the experimental results obtained after removing these modules. Precisely, the fourth, fifth, and sixth columns of Figure 9 correspond to the generated results when LST, CCPC, and MCCS are removed, respectively. It can be observed that the absence of LST significantly degrades the quality of the generated images. A possible reason is that LST effectively translates the output of the text encoder into features that the stable diffusion (Rombach et al. 2022) understands, reducing the performance degradation caused by communication gaps. Meanwhile, the lack of CCPC may cause the model to generate content that is more similar to the original image, supporting the notion that predictive coding is necessary for a task that requires generating content that has not yet occurred. Additionally, the absence of MCCS leads to a higher likelihood of semantic errors in the generated content, which is reasonable considering that MCCS provides supervision for semantic understanding.

#### Conclusion

In this study, we rethink image content generation and propose the task of causal image content generation. To support the task, we propose a dataset of VQAI based on the Tom and Jerry cartoon series. Furthermore, we analyze the challenges of the task and propose an LGD approach, which is experimentally demonstrated in this paper. Finally, we further observe the experimental results of this task on several interesting aspects and analyze some of its potential and drawbacks.

#### Acknowledgments

This work was supported by the National Key R&D Program of China (No. 2021ZD0113000) and the Innovative Development Joint Fund Key Projects of Shandong NSF (ZR2023LZH003).

## References

Aitchison, L.; and Lengyel, M. 2017. With or without you: predictive coding and Bayesian inference in the brain. *Current opinion in neurobiology*, 46: 219–227.

Antol, S.; Agrawal, A.; Lu, J.; Mitchell, M.; Batra, D.; Zitnick, C. L.; and Parikh, D. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, 2425–2433.

Balaji, Y.; Nah, S.; Huang, X.; Vahdat, A.; Song, J.; Kreis, K.; Aittala, M.; Aila, T.; Laine, S.; Catanzaro, B.; et al. 2022. ediffi: Text-to-image diffusion models with an ensemble of expert denoisers. *arXiv preprint arXiv:2211.01324*.

Bang, Y.; Cahyawijaya, S.; Lee, N.; Dai, W.; Su, D.; Wilie, B.; Lovenia, H.; Ji, Z.; Yu, T.; Chung, W.; et al. 2023. A multitask, multilingual, multimodal evaluation of chatgpt on reasoning, hallucination, and interactivity. *arXiv preprint arXiv:2302.04023*.

Brooks, T.; Holynski, A.; and Efros, A. A. 2022. Instructpix2pix: Learning to follow image editing instructions. *arXiv preprint arXiv:2211.09800*.

Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877– 1901.

Cao, Y.; Li, S.; Liu, Y.; Yan, Z.; Dai, Y.; Yu, P. S.; and Sun, L. 2023. A comprehensive survey of ai-generated content (aigc): A history of generative ai from gan to chatgpt. *arXiv preprint arXiv:2303.04226*.

Chung, H. W.; Hou, L.; Longpre, S.; Zoph, B.; Tay, Y.; Fedus, W.; Li, E.; Wang, X.; Dehghani, M.; Brahma, S.; Webson, A.; Gu, S. S.; Dai, Z.; Suzgun, M.; Chen, X.; Chowdhery, A.; Narang, S.; Mishra, G.; Yu, A.; Zhao, V.; Huang, Y.; Dai, A.; Yu, H.; Petrov, S.; Chi, E. H.; Dean, J.; Devlin, J.; Roberts, A.; Zhou, D.; Le, Q. V.; and Wei, J. 2022. Scaling Instruction-Finetuned Language Models.

Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34: 8780–8794.

Fu, R.; Zhan, X.; Chen, Y.; Ritchie, D.; and Sridhar, S. 2022. Shapecrafter: A recursive text-conditioned 3d shape generation model. *arXiv preprint arXiv:2207.09446*.

Gal, R.; Patashnik, O.; Maron, H.; Bermano, A. H.; Chechik, G.; and Cohen-Or, D. 2022. StyleGAN-NADA: CLIP-guided domain adaptation of image generators. *ACM Transactions on Graphics (TOG)*, 41(4): 1–13.

Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2020. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144.

Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33: 6840–6851.

Huang, Y.; and Rao, R. P. 2011. Predictive coding. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2(5): 580–593.

IDEA-CCNL. 2021. Fengshenbang-LM. https://github. com/IDEA-CCNL/Fengshenbang-LM.

Jahan, T.; Guan, Y.; and Van Kaick, O. 2021. Semantics-Guided Latent Space Exploration for Shape Generation. In *Computer Graphics Forum*, volume 40, 115–126. Wiley Online Library.

Kawar, B.; Zada, S.; Lang, O.; Tov, O.; Chang, H.; Dekel, T.; Mosseri, I.; and Irani, M. 2022. Imagic: Text-based real image editing with diffusion models. *arXiv preprint arXiv:2210.09276*.

Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Kingma, D. P.; and Welling, M. 2013. Auto-Encoding Variational Bayes. .

Kojima, T.; Gu, S. S.; Reid, M.; Matsuo, Y.; and Iwasawa, Y. 2022. Large language models are zero-shot reasoners. *arXiv* preprint arXiv:2205.11916.

Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023a. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*.

Li, X.; Fan, B.; Zhang, R.; Jin, L.; Wang, D.; Guo, Z.; Zhao, Y.; and Li, R. 2023b. Image Content Generation with Causal Reasoning. *arXiv preprint arXiv:2312.07132*.

Liu, Z.; Wang, Y.; Qi, X.; and Fu, C.-W. 2022. Towards implicit text-guided 3d shape generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17896–17906.

Maharana, A.; Hannan, D.; and Bansal, M. 2022. Storydalle: Adapting pretrained text-to-image transformers for story continuation. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXVII,* 70–87. Springer.

Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1): 99–106.

Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.

OpenAI. 2023. GPT-4 Technical Report. .

Pan, X.; Qin, P.; Li, Y.; Xue, H.; and Chen, W. 2022. Synthesizing Coherent Story with Auto-Regressive Latent Diffusion Models. *arXiv preprint arXiv:2211.10950*.

Qian, Y.; Fan, Y.; Hu, W.; and Soong, F. K. 2014. On the training aspects of deep neural network (DNN) for parametric TTS synthesis. In 2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 3829–3833. IEEE.

Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR. Radford, A.; Narasimhan, K.; Salimans, T.; Sutskever, I.; et al. 2018. Improving language understanding by generative pre-training.

Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.

Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1): 5485–5551.

Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*.

Ramesh, A.; Pavlov, M.; Goh, G.; Gray, S.; Voss, C.; Radford, A.; Chen, M.; and Sutskever, I. 2021. Zero-shot text-toimage generation. In *International Conference on Machine Learning*, 8821–8831. PMLR.

Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10684– 10695.

Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E. L.; Ghasemipour, K.; Gontijo Lopes, R.; Karagol Ayan, B.; Salimans, T.; et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35: 36479–36494.

Shum, K.; Diao, S.; and Zhang, T. 2023. Automatic Prompt Augmentation and Selection with Chain-of-Thought from Labeled Data. *arXiv preprint arXiv:2302.12822*.

Smiley, T. J. 1973. What is a syllogism? *Journal of philosophical logic*, 136–154.

Van Den Oord, A.; Kalchbrenner, N.; and Kavukcuoglu, K. 2016. Pixel recurrent neural networks. In *International conference on machine learning*, 1747–1756. PMLR.

Vinyals, O.; Toshev, A.; Bengio, S.; and Erhan, D. 2015. Show and tell: A neural image caption generator. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3156–3164.

Wang, J.; Zhang, Y.; Zhang, L.; Yang, P.; Gao, X.; Wu, Z.; Dong, X.; He, J.; Zhuo, J.; Yang, Q.; Huang, Y.; Li, X.; Wu, Y.; Lu, J.; Zhu, X.; Chen, W.; Han, T.; Pan, K.; Wang, R.; Wang, H.; Wu, X.; Zeng, Z.; Chen, C.; Gan, R.; and Zhang, J. 2022. Fengshenbang 1.0: Being the Foundation of Chinese Cognitive Intelligence. *CoRR*, abs/2209.02970.

Wang, L.; Xu, W.; Lan, Y.; Hu, Z.; Lan, Y.; Lee, R. K.-W.; and Lim, E.-P. 2023. Plan-and-Solve Prompting: Improving Zero-Shot Chain-of-Thought Reasoning by Large Language Models. *arXiv preprint arXiv:2305.04091*.

Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Chi, E.; Le, Q.; and Zhou, D. 2022. Chain of thought prompting elicits reasoning in large language models. *arXiv preprint arXiv:2201.11903*.

Yu, J.; Xu, Y.; Koh, J. Y.; Luong, T.; Baid, G.; Wang, Z.; Vasudevan, V.; Ku, A.; Yang, Y.; Ayan, B. K.; et al. 2022. Scaling autoregressive models for content-rich text-to-image generation. *arXiv preprint arXiv:2206.10789*.

Ze, H.; Senior, A.; and Schuster, M. 2013. Statistical parametric speech synthesis using deep neural networks. In 2013 *ieee international conference on acoustics, speech and signal processing*, 7962–7966. IEEE.

Zen, H.; and Sak, H. 2015. Unidirectional long shortterm memory recurrent neural network with recurrent output layer for low-latency speech synthesis. In 2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 4470–4474. IEEE.

Zhang, C.; Zhang, C.; Li, C.; Qiao, Y.; Zheng, S.; Dam, S. K.; Zhang, M.; Kim, J. U.; Kim, S. T.; Choi, J.; et al. 2023a. One small step for generative ai, one giant leap for agi: A complete survey on chatgpt in aigc era. *arXiv preprint arXiv:2304.06488*.

Zhang, C.; Zhang, C.; Zheng, S.; Qiao, Y.; Li, C.; Zhang, M.; Dam, S. K.; Thwal, C. M.; Tun, Y. L.; Huy, L. L.; et al. 2023b. A Complete Survey on Generative AI (AIGC): Is ChatGPT from GPT-4 to GPT-5 All You Need? *arXiv* preprint arXiv:2303.11717.

Zhang, Z.; Zhang, A.; Li, M.; and Smola, A. 2022. Automatic chain of thought prompting in large language models. *arXiv preprint arXiv:2210.03493*.

Zhang, Z.; Zhang, A.; Li, M.; Zhao, H.; Karypis, G.; and Smola, A. 2023c. Multimodal chain-of-thought reasoning in language models. *arXiv preprint arXiv:2302.00923*.