# Making Multimodal Generation Easier: When Diffusion Models Meet LLMs

**Anonymous ACL submission** 

## Abstract

We present EasyGen, an efficient model designed to enhance multimodal understanding and generation by harnessing the capabilities 004 of diffusion models and large language models (LLMs). Unlike existing multimodal models that predominately depend on encoders like 007 CLIP or ImageBind and need ample amounts of training data to bridge the gap between modalities, EasyGen leverages a bidirectional conditional diffusion model named BiDiffuser, 011 which promotes more efficient interactions between modalities. EasyGen handles image-totext generation by integrating BiDiffuser and an LLM via a simple projection layer. Unlike 015 most existing multimodal models that are limited to generating text responses, EasyGen can 017 also facilitate text-to-image generation by utilizing the LLM to create textual descriptions, which can be interpreted by BiDiffuser to gener-019 ate appropriate visual responses. Furthermore, EasyGen can be effortlessly integrated into existing advanced multimodal LLMs like LLaVA to improve their performance. Extensive quantitative and qualitative experiments demonstrate the effectiveness of EasyGen, whose training can be easily achieved in a lab setting.

## 1 Introduction

027

037

041

In recent years, remarkable progress has been made in the field of artificial intelligence generated content (AIGC), notably in technologies like large language models (LLMs) (Chiang et al., 2023; Touvron et al., 2023; Brown et al., 2020; Chowdhery et al., 2022; Zeng et al., 2022) for text generation and diffusion models (Rombach et al., 2022; Nichol et al., 2022; Saharia et al., 2022) for visual generation. These breakthroughs have paved the way for the development of large-scale multimodal generative models, sparking a recent trend of incorporating extra visual modules into LLMs. Collaborative models, such as Visual ChatGPT (Wu et al., 2023a) and MM-REACT (Yang et al., 2023), strategically use externally pre-trained tools to translate visual information into text descriptions and feed the data into LLMs. However, they are exclusively dependent on pre-trained tools for inference. Contrarily, end-to-end trained models including the BLIP series (Li et al., 2023b), LLaVA series (Liu et al., 2023b,a), MiniGPT-4 (Zhu et al., 2023), and mPLUG-Owl (Ye et al., 2023) focus on mapping image information to the text space of LLMs, enabling LLMs to comprehend visual inputs. 042

043

044

047

048

054

056

060

061

062

063

064

065

066

067

068

069

070

071

072

073

074

078

079

081

Existing end-to-end models are also not without limitations. First, most of these multimodal models rely on either CLIP (Radford et al., 2021) or Image-Bind (Girdhar et al., 2023) as their image encoder. While these encoders excel in learning unified representations that encompass both text and images, they face challenges when it comes to transforming between different modalities. This predicament makes current vision-language models relying heavily on sizable data sets to align CLIP/Bindencoded images with the language model, due to the disparity between different modalities. Furthermore, most of existing multimodal models lack the ability of generating multimodal responses, as they are primarily designed to understand multimodal content and often fall short in generating content beyond text. A recent work Emu (Sun et al., 2023) takes a unified approach to generate visual or textual tokens, but it also heavily relies on vast quantities of training data.

In this work, we propose EasyGen, a model that facilitates multimodal generation by harnessing the strengths of diffusion models and LLMs. Diffusion models with multiple skills, such as Versatile Diffusion (Xu et al., 2023) and UniDiffuser (Bao et al., 2023b), have exhibited remarkable capabilities in accurately capturing multimodal distributions. Uni-Diffuser, in particular, shows great potential in accommodating multimodal interactions. It views both image and text as sequential token streams for diffusion calculations, making it well-suited for



Figure 1: Our model EasyGen can understand multimodal inputs and generate multimodal responses, as illustrated by model-generated speech bubbles in grey color, which include both text and images.

the text representations in LLMs that are typically structured as sequences. However, since UniDiffuser aims to fit all conditional distributions (including those conditioned on noisy inputs) into one model, it is less effective on particular tasks such as conditional generation based on noise-free inputs. To address this limitation, we finetune UniDiffuser with a specific focus on the targeted image-to-text and text-to-image tasks. The finetuned model, referred to as BiDiffuser, forms a core component of EasyGen for text and image generation.

083

087

089

100

101

102

103

104

105

106

108

BiDiffuser is able to convert image data into a textual format, which simplifies the process of synchronizing its embedding space with that of an LLM for semantic comprehension and reasoning. As illustrated in Figure 2, we bridge BiDiffuser and the LLM using a simple projection layer, which can be trained efficiently with a small amount of data for image-to-text tasks such as image captioning and visual question answering. Alternatively, the LLM can be utilized to generate detailed descriptions and cues derived from text contexts like dialogues, which can aid BiDiffuser in generating accurate visual responses, as illustrated in Figure 2.



Figure 2: Overview of EasyGen.

Figure 1 demonstrates the capability of EasyGen in processing multimodal inputs and generating the

appropriate multimodal responses. Notably, Easy-Gen achieves competitive performance compared to state-of-the-art models with less training data and training time. Without employing parameterefficient fine-tuning techniques like LoRa (Hu et al., 2021), EasyGen only requires about 120 A100 (80G) GPU hours during the pre-training process (for training BiDiffuser) and 20/72 A100 (80G) GPU hours during the alignment process for FlanT5XL/Vicuna-7B (see Table 10 in Appendix H). Furthermore, EasyGen can be conveniently incorporated into contemporary advanced multimodal LLMs such as LLaVA to improve their performance, as depicted in Figure 6 in Appendix G. 109

110

111

112

113

114

115

116

117

118

119

120

121

122

123

124

125

126

127

128

129

130

131

132

133

134

# 2 Basics of Diffusion Models

**Unconditional Generation** Given a data sample taken from a real data distribution  $\mathbf{x}_0 \sim q(\mathbf{x}_0)$ , diffusion models (Sohl-Dickstein et al., 2015; Ho et al., 2020) first destruct the data by constructing a Markov forward process and gradually injecting noise to the data:

$$q(\mathbf{x}_{1:T}|\mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{x}_t|\mathbf{x}_{t-1}), \qquad (1)$$

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I}), \quad (2)$$

where  $\beta_t \in (0, 1)$  is the variance added at diffusion step t. Then, they learn to reverse the process:

$$p(\mathbf{x}_{0:T}) = p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t), \quad (3)$$

$$p_{\theta}(\mathbf{x}_{t-1}|\mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \mu_t(\mathbf{x}_t, t), \sigma_t^2 \mathbf{I}), \quad (4)$$

where  $p(\mathbf{x}_T) = \mathcal{N}(\mathbf{x}_T; 0, \mathbf{I})$  is the standard Gaussian distribution and  $\mu_t(\cdot)$  is the parameterization of the predicted mean. Diffusion models are trained 139

185

186

209

210

211

212

213

214

215

216

217

218

219

220

221

222

223

224

225

226

227

229

230

231

to maximize the marginal likelihood of the data 140  $\mathbb{E}[\log p_{\theta}(\mathbf{x}_0)]$ , and the canonical objective is the variational lower bound of  $\log p_{\theta}(\mathbf{x}_0)$ . Denoising diffusion probabilistic models (Ho et al., 2020) gen-143 erate samples  $\mathbf{x}_t \sim q(\mathbf{x}_t | \mathbf{x}_0)$  by injecting noise  $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I})$  to the data  $\mathbf{x}_0$ , and train a network  $\epsilon_{\theta}(\cdot)$  to predict the added noise  $\epsilon$  using a standard mean squared error loss:

141

142

144

145

146

147

148

150

151

152

153

154

155

156

157

158

159

162

163

164

165

166

168

171

172

173

174

175

177

178

179

180

181

182

184

$$\mathcal{L} := \mathbb{E}_{\mathbf{x}_0, \boldsymbol{\epsilon}, t} [\|\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_{\theta}(\mathbf{x}_t, t)\|^2].$$
(5)

Note that  $\mu_t(\mathbf{x}_t, t)$  can be derived from  $\epsilon_{\theta}(\mathbf{x}_t, t)$ . 149

Conditional Generation For conditional generation, a paired data  $(\mathbf{x}_0, \mathbf{y}_0)$  is given, and the aim is to model the conditional data distribution  $q(\mathbf{x}_0|\mathbf{y}_0)$ , where  $\mathbf{y}_0$  can be image class or text prompt. Conditional generation includes classifier guidance (Dhariwal and Nichol, 2021) and classifier-free guidance (Ho and Salimans, 2021). Classifier guidance requires training an extra classifier on noisy data at inference time to improve the sample quality. For classifier-free guidance, no classifier needs to be trained. The denosing network  $\epsilon_{\theta}(\mathbf{x}_t | \mathbf{y}_0)$  simply conditions on the information encoded in  $y_0$ . At inference time, with a guidance scale s, the modified score estimate is further in the direction of  $\epsilon_{\theta}(\mathbf{x}_t | \mathbf{y}_0)$  and away from the unconditional model  $\epsilon_{\theta}(\mathbf{x}_t | \emptyset)$  ( $\emptyset$  is a null token) as follows:

$$\hat{\boldsymbol{\epsilon}}_{\theta}(\mathbf{x}_t | \mathbf{y}_0) = \boldsymbol{\epsilon}_{\theta}(\mathbf{x}_t | \boldsymbol{\emptyset}) + s \cdot (\boldsymbol{\epsilon}_{\theta}(\mathbf{x}_t | \mathbf{y}_0) - \boldsymbol{\epsilon}_{\theta}(\mathbf{x}_t | \boldsymbol{\emptyset}))$$
(6)

### 3 **EasyGen: Easy Multimodal Generation** with a Bidirectional Conditional **Diffusion Model and LLMs**

We propose EasyGen, a model capable of processing multimodal inputs and generating multimodal outputs. It achieves easy multimodal generation by leveraging a bidirectional conditional diffusion model to effectively bridge the gap between different modalities and an LLM to comprehend multimodal tasks and produce textual responses containing cues for multimodal message creation. In the subsequent section, we outline the multimodal generation process of EasyGen.

#### 3.1 **BiDiffuser: A Bidirectional conditional Diffusion Model**

Since the text space of LLMs is discrete, to minimize the disparity between the output of a diffusion

model and the input of LLMs, we leverage Unidiffuser (Bao et al., 2023b), a unified diffusion model capable of transforming images into the discrete text space.

During the training process, UniDiffuser injects noise  $\epsilon^x$  and  $\epsilon^y$  to a set of paired image-text data  $(\mathbf{x}_0, \mathbf{y}_0)$  and generates noisy data  $\mathbf{x}_{t^x}$  and  $\mathbf{y}_{t^y}$ , where  $0 \leq t^x, t^y \leq T$  represent two individual timesteps (perturbation levels). It then trains a joint noise prediction network  $\epsilon_{\theta}(\mathbf{x}_{t^x}, \mathbf{y}_{t^y}, t^x, t^y)$  to predict the noise  $\epsilon^x$  and  $\epsilon^y$  by minimizing the mean squared error loss:

$$\mathbb{E}_{\boldsymbol{\epsilon}^{x},\boldsymbol{\epsilon}^{y},\mathbf{x}_{0},\mathbf{y}_{0}}[\|[\boldsymbol{\epsilon}^{x},\boldsymbol{\epsilon}^{y}]-\boldsymbol{\epsilon}_{\theta}(\mathbf{x}_{t^{x}},\mathbf{y}_{t^{y}},t^{x},t^{y})\|^{2}],\tag{7}$$

where the output of  $\epsilon_{\theta}$  is the concatenation of the estimated noise  $\epsilon_{\theta}^{x}$  and  $\epsilon_{\theta}^{y}$ , i.e.,  $\epsilon_{\theta} = [\epsilon_{\theta}^{x}, \epsilon_{\theta}^{y}]$ .

By predicting  $\epsilon_{\theta}(\mathbf{x}_{t^x}, \mathbf{y}_{t^y}, t^x, t^y)$  for any  $t^x$  and  $t^y$ , UniDiffuser learns all distributions related to  $(\mathbf{x}_0, \mathbf{y}_0)$  simultaneously. This includes all conditional distributions:  $q(\mathbf{x}_0|\mathbf{y}_0)$  for text-to-image generation,  $q(\mathbf{y}_0|\mathbf{x}_0)$  for image-to-text generation, and those conditioned on noisy input, i.e.,  $q(\mathbf{x}_0|\mathbf{y}_{t^y})$  and  $q(\mathbf{y}_0|\mathbf{x}_{t^x})$ , for  $0 < t^x, t^y \leq T$ . Learning a conditional distribution  $q(\mathbf{x}_0|\mathbf{y}_{t^y})$  or  $q(\mathbf{y}_0|\mathbf{x}_{t^x})$  can be seen as learning a distinct task. From a multitask learning perspective, due to limited network capacity, learning many tasks simultaneously (i.e., fitting all distributions to a single network) may result in task competition or task con*flict*, ultimately leading to suboptimal performance in particular tasks such as  $q(\mathbf{x}_0|\mathbf{y}_0)$  and  $q(\mathbf{y}_0|\mathbf{x}_0)$ .

To resolve this issue and enhance the performance of both image-to-text and text-to-image generation tasks, we finetune UniDiffuser with exclusive emphasis on the two tasks:

$$\mathcal{L} = \mathbb{E}_{\boldsymbol{\epsilon}^{x}, \boldsymbol{\epsilon}^{y}, \mathbf{x}_{0}, \mathbf{y}_{0}} [\|\boldsymbol{\epsilon}^{x} - \boldsymbol{\epsilon}_{\theta}^{x}(\mathbf{x}_{t^{x}}, \mathbf{y}_{0}, t^{x}, 0)\|^{2} + \alpha \|\boldsymbol{\epsilon}^{y} - \boldsymbol{\epsilon}_{\theta}^{y}(\mathbf{x}_{0}, \mathbf{y}_{t^{y}}, 0, t^{y})\|^{2}].$$
(8)

where  $\alpha$  is a hyperparameter to balance the learning paces of the two tasks. As depicted in Figure 3, our training objective entails predicting the text  $\mathbf{y}_0$  based on the input image  $\mathbf{x}_0$  and vice versa, where the input conditions for the model are noisefree. We employ classifier-free guidance. During training, we estimate the noise injected to the image (i.e.,  $\epsilon_{\theta}^{x}(\mathbf{x}_{t^{x}}, \mathbf{y}_{0}, t^{x}, 0)$ ) conditioned on the noise-free text  $\mathbf{y}_0$  and the noise to the text (i.e.,  $\epsilon^y_{\theta}(\mathbf{x}_0, \mathbf{y}_{t^y}, 0, t^y))$  given the noise-free image  $\mathbf{x}_0$ . During inference, with a guidance scale  $s \ge 0$ , we use the modified prediction  $\hat{\epsilon}_{\theta}$  to guide towards the



Figure 3: The training process of BiDiffuser involves finetuning UniDiffuser (Bao et al., 2023b) with a joint objective of image-to-text and text-to-image tasks.

condition:

$$\hat{\boldsymbol{\epsilon}}_{\theta}^{x}(\mathbf{x}_{t^{x}}, \mathbf{y}_{0}, t^{x}, 0) = \boldsymbol{\epsilon}_{\theta}^{x}(\mathbf{x}_{t^{x}}, \boldsymbol{\epsilon}^{y}, t^{x}, T) + s \cdot (\boldsymbol{\epsilon}_{\theta}^{x}(\mathbf{x}_{t^{x}}, \mathbf{y}_{0}, t^{x}, 0) - \boldsymbol{\epsilon}_{\theta}^{x}(\mathbf{x}_{t^{x}}, \boldsymbol{\epsilon}^{y}, t^{x}, T)), \\ \hat{\boldsymbol{\epsilon}}_{\theta}^{y}(\mathbf{x}_{0}, \mathbf{y}_{t^{y}}, 0, t^{y}) = \boldsymbol{\epsilon}_{\theta}^{y}(\boldsymbol{\epsilon}^{x}, \mathbf{y}_{t^{y}}, T, t^{y}) + s \cdot (\boldsymbol{\epsilon}_{\theta}^{y}(\mathbf{x}_{0}, \mathbf{y}_{t^{y}}, 0, t^{y}) - \boldsymbol{\epsilon}_{\theta}^{y}(\boldsymbol{\epsilon}^{x}, \mathbf{y}_{t^{y}}, T, t^{y})).$$
(9)

where  $\epsilon_{\theta}^{x}(\mathbf{x}_{t^{x}}, \epsilon^{y}, t^{x}, T)$  ( $t^{y} = T$  and  $\mathbf{y}_{T} = \epsilon^{y}$ ) and  $\epsilon_{\theta}^{y}(\epsilon^{x}, \mathbf{y}_{t^{y}}, T, t^{y})$  ( $t^{x} = T$  and  $\mathbf{x}_{T} = \epsilon^{x}$ ) represent the unconditional models when T is sufficiently large. We name the finetuned model "BiDiffuser", signifying its specialized ability in bidirectional conditional generation.

### 3.2 Image-to-Text Generation

BiDiffuser can convert images into vectors in the text space, facilitating alignment with the vector space of LLMs. In the following, we show how BiDiffuser can be integrated with LLMs to perform image-to-text generation tasks such as image captioning and visual question answering (VQA).

### 3.2.1 Aligning BiDiffuser with LLMs

We connect BiDiffuser and LLMs via a simple projection layer, which maps text embeddings obtained from the output of the diffusion model to the embedding space of LLMs. As shown in Figure 4, the alignment can take place either prior to the LLM (referred to as Pre-Align manner) or between its encoder and decoder components (referred to as Mid-Align manner).

Pre-Align Manner As shown in Figure 4a, the
projection layer is placed before the LLM to map
the output of BiDiffuser (image representations)
to the text embedding space of the LLM. The text
embedding of the input image is then concatenated
with the embeddings of the textual instructions and

fed to the LLM for decoding. To synchronize the text space of BiDiffuser with that of the LLM, we propose to use the image-grounded text generation (ITG) objective to drive the model to generate texts based on the input image by computing the autoregressive loss:

262

263

264

265

266

270

271

272

273

274

275

276

277

279

280

281

284

285

$$\mathcal{L}_{\text{ITG}} = -\frac{1}{L} \sum_{l=1}^{L} \log p_{\theta}(w_l^g | w_{< l}^g, I, T_I), \quad (10)$$

where  $w^g = (w_1^g, ..., w_L^g)$  represents the groundtruth caption of image *I* with length *L*, *T<sub>I</sub>* is the text instruction, and  $\theta$  denotes the model parameters, which include the parameters of the projection layer and the LLM.

**Mid-Align Manner** As shown in Figure 4b, the projection layer is placed between the LLM's encoder and decoder, aiming to map the output of BiDiffuser to the embedding space of the text that is encoded by the LLM's encoder. Particularly, we argue that the output of BiDiffuser, once mapped by the projection layer and denoted as  $d_{diff}$ , should align with the image caption that is encoded by the LLM's encoder as  $d_{llm}$ . Therefore, to accurately learn the alignment between the image and text representations, in addition to the ITG loss in Eq. 10, we also employ an image-text distance minimization (ITDM) loss:

$$\mathcal{L}_{\text{ITDM}} = \frac{1}{N} \sum_{i=1}^{N} \|\mathbf{d}_{\text{diff}} - \mathbf{d}_{\text{llm}}\|_{2}^{2}, \qquad (11)$$
$$\mathcal{L}_{\text{mid}} = \mathcal{L}_{\text{ITG}} + \mathcal{L}_{\text{ITM}}.$$

where N is the batch size, and  $\mathcal{L}_{mid}$  is the overall288loss. In this manner, the model parameters  $\theta$  only289include the parameters of the projection layer.290

240

241

243

245

247



Figure 4: Two different ways of aligning BiDiffuser with LLMs.

After aligning BiDiffuser with LLMs, EasyGen gains the capability of zero-shot image-to-text generation, which includes tasks such as image captioning and VQA.

## 3.2.2 Instruction-Tuning LLMs

291

293

294

297

298

305

307

312

313

When aligning BiDiffuser with an LLM, we perform instruction-tuning on the LLM to equip it with the capability of understanding multimodal tasks. We designed different instructions for different LLMs, as shown in Table 1. General instruction template is denoted as follows:

USER: <Img><image></Img> + Instruction. Assistant: <answer>.

For the <image> placeholder, we substitute it with the output of BiDiffuser. To avoid over fitting to the specific task and counter the model's inclination to generate excessively short outputs, we have devised specific instructions (blue texts in Table 9), which enable the LLM to produce concise responses when necessary. For different tasks, the distinct instruction templates are as outlined in Appendix G.

## 3.3 Text-to-Image Response Generation

Most of existing multimodal models, including the BLIP series and LLaVA series are unable to provide a multimodal response as they are primarily designed to generate only textual outputs. On the other hand, Emu (Sun et al., 2023) takes a unified approach to predict the subsequent visual or textual token in an auto-regressive manner, but it is heavily reliant on vast quantities of training data. Contrary to the limitations of these existing models, Easy-Gen, by leveraging the bidirectional generation ca-

Different Inst	truction Templates:
Caption Generation Response	Provide a brief description of the given im- age. Assistant: Buses parked with a snow mountain view be- hind them.
VQA Response	What numbers are displayed on the front of the bus on the right? <b>Short answer: Assistant:</b> 6044
Multimodal Dialogue Response	What might be the purpose of the buses in this location? <b>Assistant:</b> It is plausible that they are in this location for multiple reasons. Some possible reasons might be: 1)2)

Table 1: Examples of different instructions with different output formats. We use bold fonts to indicate different instructions.

pability of BiDiffuser and the inference capability of LLMs, can produce accurate and high-quality visual response with ease.

To tackle multimodal response generation tasks such as PhotoChat (Zang et al., 2021), we first finetune the LLM to generate detailed image captions based on dialogue context. Then, we employ BiDiffuser to create the corresponding images with the produced captions. Specifically, we replace the image featured in the dialogue with its corresponding descriptive caption, encapsulating it with taskspecific tokens <Img>,</Img> and constructing the following instruction templates:

USER: Dialog history + <Img><caption></Img> + Dialog history. Assistant: <response>.

USER: Dialog history. Assistant: <response> +



Figure 5: Text-to-image generation by EasyGen. (Bottom) LLM generates response and description of the image. (Top) BiDiffuser takes the description as input and generates images.

When <caption> appears in response, it represents the generated description of the image. Training with the instruction data enables our model to not only produce text responses but also perform image intent classification and generate image captions that BiDiffuser can interpret.

## 4 Experiments

## 4.1 Experimental Setup

We initialize the encoder-decoder LLM using pretrained weights from FlanT5XL or decoder-only LLM from Vicuna-7B, along with the utilization of the diffusion module from BiDiffuser. During the alignment process, we maintain the frozen state of the BiDiffuser. The statistics of the datasets for pre-training, alignment and instruction-tuning can be found in Appendix C. For the image captioning task, EasyGen is evaluated on both the MS-COCO (Lin et al., 2014) Karpathy test set and the NoCaps (Agrawal et al., 2019) validation set. For the VQA task, our method is evaluated on OK-VQA (Marino et al., 2019) validation set and GQA (Hudson and Manning, 2019) test-dev set.

To adapt the model for multimodal dialogue generation, we fine-tune the LLM and projection layer on the PhotoChat dataset. We incorporate photosharing activities into the dialogue context by generating <Img><caption></Img>, and utilize crossentropy loss exclusively for fine-tuning the multimodal generation task. Given the limited expressiveness of image descriptions in the PhotoChat dataset, as evidenced by Figure 7 in Appendix J, we regenerate image annotations in a text format similar to that used in MS-COCO.

## 4.2 Evaluation

We evaluate EasyGen on various vision-language tasks including image captioning (MS-COCO (Lin et al., 2014), NoCaps (Agrawal et al., 2019)), visual question answering (OK-VQA (Marino et al., 2019), GQA (Hudson and Manning, 2019)), and multimodal dialog generation (PhotoChat (Zang et al., 2021)). We use BLIP (Li et al., 2022), Flamingo (Alayrac et al., 2022), BLIP-2 (Li et al., 2023b), InstructBlip (Dai et al., 2023), MiniGPT-4 (Zhu et al., 2023), and LLaVA (Liu et al., 2023b) as baselines for image-to-text tasks, and Maria (Liang et al., 2021) and Divter (Sun et al., 2021) as baselines for the multimodal response generation task. See details in Appendix D and F. 374

375

376

377

378

380

381

385

386

387

389

390

391

392

393

394

395

396

397

398

399

400

401

402

403

404

405

406

407

408

409

410

411

412

413

414

415

416

417

418

419

420

421

422

423

## 4.3 Overall Results

Table 2 lists the evaluation results for each baseline and our models on MS-COCO and VQA (zeroshot) datasets. EasyGen outperforms most of the baseline models on both the COCO test set and No-Caps validation set (zero-shot transfer). Although EasyGen is only pre-trained on a small dataset MS-COCO, its performance on the image captioning generation task is comparable to models (e.g., BLIP-2) pre-trained on a large dataset. This indicates that EasyGen can effectively combine the strength of diffusion module and LLM to generate smooth and informative captions. GPT scores do not vary significantly because the captions produced by the models in the image-captioning task tend to be quite alike. For the OK-VQA and GQA dataset, the performance of EasyGen is improved compared with other models of a similar scale. For example, BLIP-2 adopts the task-special decoding method and achieves 39.4% accuracy on OK-VQA validation set, while ours can get 45.2% even with a simple decoding method, i.e., greedy search.

Table 3 lists the evaluation results on the PhotoChat dataset. The results of Divter are cited from (Sun et al., 2021). We fine-tune Maria on PhotoChat dataset only for the response generation task. Since our EasyGen model can generate response and image description simultaneously, the response and description generation task has a similar PPL. Compared with other models, our method has clear advantages in the performance of PPL, indicating that our model demonstrates strong performance on text generation tasks. Besides, we find that the image descriptions in the PhotoChat dataset are too concise to convey the information of images.

341

345

347

353

354

355

361

363

371

Model	Dataset Size		NoCaps (val)		COCO (Karpathy)		OK-VQA	GQA
	PT	IT	CIDEr	SPICE	BLEU@4	CIDEr	Accuracy	Accuracy
BLIP (Li et al., 2022)	129M	-	113.2	14.8	40.4	136.7	-	-
Flamingo (Alayrac et al., 2022)	1.8B	-	-	-	-	138.1	50.6	-
BLIP-2 OPT-6.7B (Li et al., 2023b)	129M	-	121.0	15.3	43.5	145.2	36.4	36.4
BLIP-2 FlanT5XL (Li et al., 2023b)	129M	-	121.6	15.8	42.4	144.5	39.4	44.4
InstructBlip 7B (Dai et al., 2023)	129M	1.2M	123.1	-	40.8	140.7	61.0*	$49.2^{\star}$
MiniGPT-4 (Zhu et al., 2023)	-	5M	42.4	-	-	-	37.5	30.8
LLaVA (Liu et al., 2023b)	558K	158K	33.1	-	7.9	30.0	54.4	41.3
EasyGen FlanT5XL EasyGen Vicuna-7B	169K 169K	90K 90K	121.2 121.8	15.5 <b>15.8</b>	<b>43.5</b> 42.4	<b>145.7</b> 144.6	41.1 45.2	37.2 <b>44.6</b>

Table 2: Automatic evaluation and GPT evaluation of our model and the baselines on various vision-language tasks. PT, IT indicate the number of samples in pretraining and instruction tuning stage respectively. The results of EasyGen on NoCaps, OK-VQA and GQA are obtained in a zero-shot setting.  $\star$  indicates that the model was trained on other VQA datasets.

Model		<b>Response Generation</b>				<b>Description Generation</b>			
	BLEU-1	BLEU-2	PPL↓	ROUGE-L	BLEU-1/2	ROUGE-L	PPL↓	FID↓	
Divter (Sun et al., 2021) Maria (Liang et al., 2021)	6.5 13.8	1.7 9.2	59.63 48.75	5.69 15.17	15.1/11.4	15.81	5.12	29.16	
EasyGen FlanT5XL + w/ generated desc. + w/o generated desc.	<b>22.3</b> 17.8	<b>18.7</b> 12.4	<b>4.32</b> 7.61	<b>17.24</b> 15.12	13.5/10.2 17.4/13.2	13.84 <b>16.71</b>	<b>4.16</b> 6.23	<b>10.30</b> 75.46	

Table 3: Automatic evaluation of our model and the baselines on the PhotoChat dataset.

494 Therefore, we used the pre-trained model from the first stage to regenerate the image description (re-425 ferred to as "w/ generated desc." in Table 3) which 426 led to a large gap towards ground-truth descrip-427 tions, resulting in lower BLEU-1/2 and ROUGE-L. 428 However, the performance of our model on BELU-429 430 1/2 and ROUGE is higher than other models on response generation tasks, indicating that introducing 431 richer image descriptions is beneficial for generat-432 ing more relevant and informative responses. 433

### 4.4 Ablation Study

434

435 In Table 4, we examine the impact of freezing/tuning BiDiffuser and the LLM. We conducted 436 ablation studies on image captioning and VQA 437 tasks. It can be observed that the frozen Mid-Align 438 method outperforms the Pre-Align method in im-439 age captioning. This shows that the ITDM loss 440 function is effective. However, the frozen Mid-441 Align method exhibits inferior performance in the 442 VQA task. We hypothesize that this is due to the 443 integration of mid-aligned target image features 444 445 with query information, and the projection layer is insensitive to instruction information. We con-446 duct instruction-tuning on Pre-Align T5 and Vi-447 cuna. Compared to models at the same scale, these 448 instruction-tuned models achieve superior results. 449

### 4.5 Fine-tuning for VQA Tasks

In order to enhance the performance of EasyGen on VQA and OCR tasks, we take into account the potential information dilution or omission that may occur when using BiDiffuser to convert images to text space. So we opt to integrate BiDiffuser with CLIP ViT-L/14 as image encoder (as shown in Figure 6 in Appendix). In this process, we freeze CLIP and BiDiffuser while fine-tuning the parameters of the LLM and projection layers. Table 5 presents the results on traditional short QA and the modern benchmark MMBench (Liu et al., 2023c). For finetuning EasyGen, we utilize train and val splits from VQAv2, Text Captions, AOK-VQA, and TextVQA datasets. Though we do not pre-align CLIP with LLM and use smaller instruction data for tuning, EasyGen's performance is better than LLaVA on VQAv2 dataset, which shows BiDiffuser can effectively help LLM understand images. Furthermore, BiDiffuser can also be plug-and-played into other MLLMs and improve their performance. We speculate that BiDiffuser provides guidance information to MLLMs, enabling them to better understand the details of CLIP encoded images.

450

451

452

453

454

455

456

457

458

459

460

461

462

463

464

465

466

467

468

469

470

471

472

473

474

475

476

477

## 5 Related Work

Multimodal Language Models. Recent research has witnessed a surge of interest in multimodal LLMs, including collaborative models

LLM	Diffusion	Alignment	NoCaps	CC	OCO(Karpatl	hy)	OK-VQA
	Model	8	CIDEr	SPICE	BLEU@4	CIDEr	Accuracy
🏶 T5	UniDiffuser	Pre-Align	62.4	18.0	26.8	90.7	33.0
Ċ T5	BiDiffuser	Pre-Align	119.1	25.5	42.6	145.1	41.1
🏶 T5	BiDiffuser	Mid-Align	121.2	25.1	43.5	145.7	31.5
<b>O</b> T5	BiDiffuser	Mid-Align	121.5	25.3	43.6	145.7	36.4
😊 Vicuna-7B 🍀 Vicuna-7B	BiDiffuser BiDiffuser	Pre-Align Pre-Align	<b>121.8</b> 119.0	24.9 24.6	42.4 40.3	144.6 140.3	<b>45.2</b> 42.7

Table 4: Ablation studied on image captioning and VQA tasks. O / \* represents we tune/freeze the weights of the LLM during the alignment process.

Model	IT	VQAv2 (test-dev)	TextVQA	MMbench (test)
MiniGPT-4 (Zhu et al., 2023) InstructBLIP Vicuna-7B (Dai et al., 2023) LLaVA-1.5 Vicuna-7B (Liu et al., 2023a) LLaVA-1.5 Vicuna-13B (Liu et al., 2023a) EasyGen Vicuna-7B w/ ViT-L LLaVA-1.5 Vicuna-7B w/ EasyGen LLaVA-1.5 Vicuna-13B w/ EasyGen	5M 1.2M 665K 665K 251K 665K 665K	- 78.5 80.0 79.4 80.2 80.5	19.4 50.1 58.2 61.3 57.9 58.8 <b>61.5</b>	23.0 33.9 65.2 67.8 63.9 66.1 <b>69.2</b>

Table 5: Automatic evaluation of our model variants and baselines on more VQA datasets and the MMBench benchmark. "w/ EasyGen" means incorporating the core components of our model into existing models as depicted in Figure 6 in Appendix G. EasyGen variants rank among the top models on the leaderboard of MMBench.

such as Visual ChatGPT (Wu et al., 2023a), MM-REACT (Yang et al., 2023), and Hugging-GPT (Shen et al., 2023), and end-to-end methods including Flamingo (Alayrac et al., 2022), Img2LLM (Guo et al., 2022), BLIP series (Li et al., 2023b; Dai et al., 2023; Li et al., 2022), BEiT series (Bao et al., 2021; Wang et al., 2022b), LLaVA (Liu et al., 2023b), mPLUG-owl (Ye et al., 2023), MiniGPT-4 (Zhu et al., 2023), Llamaadapter (Zhang et al., 2023a), Otter (Li et al., 2023a), OFA (Wang et al., 2022a), and PaLI (Chen et al., 2022). EasyGen is built upon a bidirectional conditional diffusion model, which promotes more efficient interactions between modalities.

478

479

480

481

483

484

485

486

487

488

489

490

491

Multimodal Diffusion Models. Diffusion genera-492 493 tive models (Rombach et al., 2022; Ramesh et al., 2021; Nichol et al., 2022; Ruiz et al., 2023) have 494 achieved strong results in text conditioned image 495 generation works. Specifically, Versatile Diffu-496 sion (Xu et al., 2023) employs the U-Net (Ron-497 neberger et al., 2015) architecture with a multi-498 flow design to tackle multiple modalities and tasks, 499 while UniDiffuser (Bao et al., 2023b) adopts the U-ViT (Bao et al., 2023a) framework to treat both image and text as sequential token streams for dif-503 fusion calculations. However, these models are unable to complete complex language tasks. Easy-504 Gen combines the advantages of diffusion models and LLMs and achieves competitive performance in both image-to-text and text-to-image tasks. 507

**Multimodal Response Generation**. Recent works have shown significant progress on multimodal response generation (Koh et al., 2023b; Aghajanyan et al., 2022; Zhang et al., 2023b; Wu et al., 2023b; Pan et al., 2023; Koh et al., 2023a). Divter (Sun et al., 2021) incorporates text-to-image generation into text-only dialogue response generation to produce a multimodal response. Leveraging the power of diffusion models, CoDi (Tang et al., 2023) can generate any combination of output modalities. Emu (Sun et al., 2023) takes a unified approach to predict the subsequent visual or textual token in an auto-regressive manne. In EasyGen, we efficiently combine the diffusion model and LLMs to generate multimodal outputs.

508

509

510

511

512

513

514

515

516

517

518

519

520

521

522

523

524

525

526

528

529

530

531

532

533

534

535

536

# 6 Conclusion

We have introduced EasyGen, a model that facilitates multimodal understanding and generation. In contrast to existing models that rely on encoders like CLIP or ImageBind and require significant amounts of training data to integrate different modalities, EasyGen offers a more efficient solution by employing a bidirectional diffusion model named BiDiffuser. This allows for more effective modal interactions, handling both image-to-text and text-to-image generations by the fusion of BiDiffuser and LLMs. Additionally, EasyGen can be easily integrated into existing advanced multimodal LLMs to further boost their performance.

## References

537

538

539

540

541

543

545

546

547

548

549

553

556

563 564

565

569

570

571

576

577

579

580

583

584

585

586

587

589

593

- Armen Aghajanyan, Bernie Huang, Candace Ross, Vladimir Karpukhin, Hu Xu, Naman Goyal, Dmytro Okhonko, Mandar Joshi, Gargi Ghosh, Mike Lewis, et al. 2022. Cm3: A causal masked multimodal model of the internet. *arXiv preprint arXiv:2201.07520*.
- Harsh Agrawal, Karan Desai, Yufei Wang, Xinlei Chen, Rishabh Jain, Mark Johnson, Dhruv Batra, Devi Parikh, Stefan Lee, and Peter Anderson. 2019. Nocaps: Novel object captioning at scale. In Proceedings of the IEEE/CVF international conference on computer vision, pages 8948–8957.
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736.
- Fan Bao, Shen Nie, Kaiwen Xue, Yue Cao, Chongxuan Li, Hang Su, and Jun Zhu. 2023a. All are worth words: A vit backbone for diffusion models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 22669–22679.
- Fan Bao, Shen Nie, Kaiwen Xue, Chongxuan Li, Shi Pu, Yaole Wang, Gang Yue, Yue Cao, Hang Su, and Jun Zhu. 2023b. One transformer fits all distributions in multi-modal diffusion at scale. *arXiv preprint arXiv:2303.06555*.
- Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei.
  2021. Beit: Bert pre-training of image transformers.
  In *International Conference on Learning Representations*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, et al. 2022. Pali: A jointly-scaled multilingual language-image model. In *The Eleventh International Conference on Learning Representations*.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. Vicuna: An opensource chatbot impressing gpt-4 with 90%\* chatgpt quality.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. Palm: Scaling language modeling with pathways. arXiv preprint arXiv:2204.02311.

- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. 2023. Instructblip: Towards general-purpose vision-language models with instruction tuning.
- Prafulla Dhariwal and Alexander Nichol. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794.
- Rohit Girdhar, Alaaeldin El-Nouby, Zhuang Liu, Mannat Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. 2023. Imagebind: One embedding space to bind them all. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15180–15190.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913.
- Jiaxian Guo, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Boyang Li, Dacheng Tao, and Steven CH Hoi. 2022. From images to textual prompts: Zero-shot vqa with frozen large language models. *arXiv preprint arXiv:2212.10846*.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840– 6851.
- Jonathan Ho and Tim Salimans. 2021. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. 2022. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*.
- Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2021. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*.
- Drew A Hudson and Christopher D Manning. 2019. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6700–6709.

638

639

640

641

642

643

644

645

646

647

- Jing Yu Koh, Daniel Fried, and Ruslan Salakhutdinov. 2023a. Generating images with multimodal language models. *arXiv preprint arXiv:2305.17216*.
  - Jing Yu Koh, Ruslan Salakhutdinov, and Daniel Fried. 2023b. Grounding language models to images for multimodal generation. *arXiv preprint arXiv:2301.13823*.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123:32–73.

666

670

671

675

676

677

679

697

698

702

- Bo Li, Yuanhan Zhang, Liangyu Chen, Jinghao Wang, Jingkang Yang, and Ziwei Liu. 2023a. Otter: A multi-modal model with in-context instruction tuning. *arXiv preprint arXiv:2305.03726*.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023b. Blip-2: Bootstrapping language-image pretraining with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pretraining for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR.
- Zujie Liang, Huang Hu, Can Xu, Chongyang Tao, Xiubo Geng, Yining Chen, Fan Liang, and Daxin Jiang.
  2021. Maria: A visual experience powered conversational agent. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics*, pages 5596–5611.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In Computer Vision– ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13, pages 740–755. Springer.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2023a. Improved baselines with visual instruction tuning. *arXiv preprint arXiv:2310.03744*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023b. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*.
- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. 2023c. Mmbench: Is your multi-modal model an all-around player? *arXiv preprint arXiv:2307.06281*.
- Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. 2022. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *arXiv preprint arXiv:2211.01095*.

Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. 2019. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *Proceedings of the IEEE/cvf conference on computer vision and pattern recognition*, pages 3195–3204. 703

704

706

707

709

711

712

713

714

715

716

717

718

719

720

721

722

723

724

725

726

727

728

729

730

731

732

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756

757

758

- Vishvak Murahari, Prithvijit Chattopadhyay, Dhruv Batra, Devi Parikh, and Abhishek Das. 2019. Improving generative visual dialog by answering diverse questions. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Alexander Quinn Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob Mcgrew, Ilya Sutskever, and Mark Chen. 2022. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. In *International Conference on Machine Learning*, pages 16784–16804. PMLR.
- Xichen Pan, Li Dong, Shaohan Huang, Zhiliang Peng, Wenhu Chen, and Furu Wei. 2023. Kosmos-g: Generating images in context with multimodal large language models. *arXiv preprint arXiv:2310.02992*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. Highresolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695.
- 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: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18, pages 234–241. Springer.
- Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. 2023. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22500–22510.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. 2022. Photorealistic text-to-image diffusion models with deep

866

867

868

869

870

871

815

816

817

771

760

761

- 775 776 777
- 782

787

- 793
- 794

798

804

810 811

812

- language understanding. Advances in Neural Information Processing Systems, 35:36479-36494.
- Dustin Schwenk, Apoorv Khandelwal, Christopher Clark, Kenneth Marino, and Roozbeh Mottaghi. 2022. A-okvqa: A benchmark for visual question answering using world knowledge. In European Conference on Computer Vision, pages 146-162. Springer.
- Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. 2023. Hugginggpt: Solving ai tasks with chatgpt and its friends in huggingface. arXiv preprint arXiv:2303.17580.
- Oleksii Sidorov, Ronghang Hu, Marcus Rohrbach, and Amanpreet Singh. 2020. Textcaps: a dataset for image captioning with reading comprehension. In Computer Vision-ECCV 2020: 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part II 16, pages 742-758. Springer.
- Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. 2019. Towards vqa models that can read. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In International conference on machine learning, pages 2256–2265. PMLR.
- Qingfeng Sun, Yujing Wang, Can Xu, Kai Zheng, Yaming Yang, Huang Hu, Fei Xu, Jessica Zhang, Xiubo Geng, and Daxin Jiang. 2021. Multimodal dialogue response generation. arXiv preprint arXiv:2110.08515.
- Quan Sun, Qiying Yu, Yufeng Cui, Fan Zhang, Xiaosong Zhang, Yueze Wang, Hongcheng Gao, Jingjing Liu, Tiejun Huang, and Xinlong Wang. 2023. Generative pretraining in multimodality. arXiv preprint arXiv:2307.05222.
- Zineng Tang, Ziyi Yang, Chenguang Zhu, Michael Zeng, and Mohit Bansal. 2023. Any-to-any generation via composable diffusion. arXiv preprint arXiv:2305.11846.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. arXiv preprint arXiv:2302.13971.
- Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. 2022a. Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In International Conference on Machine Learning, pages 23318-23340. PMLR.

- Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, et al. 2022b. Image as a foreign language: Beit pretraining for all vision and vision-language tasks. arXiv preprint arXiv:2208.10442.
- Chenfei Wu, Shengming Yin, Weizhen Qi, Xiaodong Wang, Zecheng Tang, and Nan Duan. 2023a. Visual chatgpt: Talking, drawing and editing with visual foundation models. arXiv preprint arXiv:2303.04671.
- Shengqiong Wu, Hao Fei, Leigang Qu, Wei Ji, and Tat-Seng Chua. 2023b. Next-gpt: Any-to-any multimodal llm. arXiv preprint arXiv:2309.05519.
- Xingqian Xu, Zhangyang Wang, Gong Zhang, Kai Wang, and Humphrey Shi. 2023. Versatile diffusion: Text, images and variations all in one diffusion model. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 7754–7765.
- Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Ehsan Azarnasab, Faisal Ahmed, Zicheng Liu, Ce Liu, Michael Zeng, and Lijuan Wang. 2023. Mmreact: Prompting chatgpt for multimodal reasoning and action. arXiv preprint arXiv:2303.11381.
- Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, et al. 2023. mplug-owl: Modularization empowers large language models with multimodality. arXiv preprint arXiv:2304.14178.
- Xiaoxue Zang, Lijuan Liu, Maria Wang, Yang Song, Hao Zhang, and Jindong Chen. 2021. Photochat: A human-human dialogue dataset with photo sharing behavior for joint image-text modeling. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pages 6142-6152.
- Aohan Zeng, Xiao Liu, Zhengxiao Du, Zihan Wang, Hanyu Lai, Ming Ding, Zhuoyi Yang, Yifan Xu, Wendi Zheng, Xiao Xia, et al. 2022. Glm-130b: An open bilingual pre-trained model. In The Eleventh International Conference on Learning Representations.
- Renrui Zhang, Jiaming Han, Aojun Zhou, Xiangfei Hu, Shilin Yan, Pan Lu, Hongsheng Li, Peng Gao, and Yu Qiao. 2023a. Llama-adapter: Efficient fine-tuning of language models with zero-init attention. arXiv preprint arXiv:2303.16199.
- Yiyuan Zhang, Kaixiong Gong, Kaipeng Zhang, Hongsheng Li, Yu Qiao, Wanli Ouyang, and Xiangyu Yue. 2023b. Meta-transformer: A unified framework for multimodal learning. arXiv preprint arXiv:2307.10802.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2023. Minigpt-4: Enhancing vision-language understanding with advanced large language models. arXiv preprint arXiv:2304.10592.

875

876

878

883

887

893

894

901

902

903

904

905

906

907

908

909

910

911

912

913

914

916

917

918

919

920

# A Limitations

This section aims to highlight the limitations of our work and provide further insights into the research in this area. Our model relies on diffusion for multimodal interaction, which means that the text-toimage and image-to-text processes may take longer. In our experiments, we tested the performance of our model on one A100(80G) GPU. During inference, using 1000 image-caption pairs, EasyGen took approximately 2.95 seconds for the caption generation task (with the diffusion module taking about 2.41 seconds) and around 4.96 seconds to generate an image. We believe it would be beneficial to explore more efficient sampling methods, such as DPM-Solver++(Lu et al., 2022), to improve the overall efficiency of EasyGen.

## **B** Ethics Statement

We adhere to the ACL Ethics Policy and have conducted our research using publicly available repositories and datasets. Our primary focus is on investigating the integration of diffusion models and LLMs for multimodal generation. Therefore, the results should be seen as AI-generated content. While we have not observed deliberate harmful content, the model has the potential to generate such content if triggered. We have taken steps to minimize this risk through fine-tuning on public datasets, but caution is still exercised. In future, we will prioritize improving downstream performance and exploring methods to enhance control over the generation process. To ensure reproducibility and support future research, we have made all resources publicly available and provided proper citations to previous research within the code.

# C Datasets

We test the effectiveness of EasyGen by experimenting on different tasks including image captioning, visual question answering (VQA), and multimodal dialogue tasks. Table 6 shows the statistics of the pre-training datasets for BiDiffuser, alignment and VQA tasks.

We use the MS-COCO (Lin et al., 2014) dataset for image captioning. Following BLIP-2 (Li et al., 2023b), we fine-tune EasyGen on MS-COCO and evaluate its performance on the Karpathy test set and the NoCaps (Agrawal et al., 2019) validation set. In MS-COCO, each image typically has five captions that convey similar meanings. The training set consists of 82,783 images with 414,113 captions, while the COCO Karpathy test set has 5,000 images and the NoCaps validation set has 4,500 images.

For multimodal dialogue, we utilize the PhotoChat (Zang et al., 2021) dataset, which is a highquality dataset consisting of 10,917 images and 12,286 dialogues. Each dialogue is associated with a user image and its corresponding text description. The dataset is divided into 10,286 training instances, 1,000 development instances, and 1,000 testing instances. Moreover, PhotoChat includes photo-sharing activities, defined as the process of creating <Img><caption></Img> in this study. Each conversation in PhotoChat is broken down and constructed into multiple samples so that each round of responses can be learned. Specifically, we regard the first three turns as the dialog context, and the subsequent turns as the prediction targets. By converting the dialogues of this dataset into the form mentioned in 3.3, we obtained 49,240 train, 4,792 dev, and 4,836 test dialogue pairs.

For the VQA task, we conduct a quantitative evaluation on both the OK-VQA (Marino et al., 2019) validation set (5,046 questions) and the GQA (Hudson and Manning, 2019) test-dev set (12,578 questions). As shown in Table 4, for the frozen LLM, following BLIP-2, we employ the length penalty in beam search to encourage short answer generation. On the contrary, for the tuned LLM, we use the VQA instructions (as shown in Table 8) to do instruction tuning during the alignment process. The data for instruction tuning is constructed by randomly selecting 5K data from VQAv2 (Goyal et al., 2017) and 5K data from Visual Dialog (Murahari et al., 2019) training set.

# **D** Baselines

We compare our proposed model with the following state-of-the-art baselines:

**BLIP** (Li et al., 2022) is a multimodal mixture of encoder-decoder. It can be used as an image-based text encoder or decoder. We use it to perform caption generation and VQA tasks.

**BLIP-2** (Li et al., 2023b) is pre-trained through bootstrapped learning from frozen visual encoder and LLMs using an efficient pre-training strategy. **Flamingo** (Alayrac et al., 2022) incorporates new cross-attention layers into Chinchilla language model (Hoffmann et al., 2022) to inject visual features, and pre-trains the new layers on billions of image-text pairs. We use it to perform caption gen923 924 925

926

927

928

929

930

931

932

933

934

935

936

937

938

939

940

941

942

943

944

945

946

947

948

921

922

962

963

964

965

966

967

968

969

Data types	Dataset	Size	BiDiffuser	Alignment	Fine-tuning
Caption	MS-COCO caption (Lin et al., 2014) Visual Genome (Krishna et al., 2017)	83K 86K	~ ~	×	××
Multimodal instruction	LLaVA dataset (Liu et al., 2023b)	80K	×	<ul> <li>✓</li> </ul>	<ul> <li>✓</li> </ul>
VQA	VQAv2 (Goyal et al., 2017) AOK-VQA (Schwenk et al., 2022)	83K 66K	× ×	×	<i>v</i> <i>v</i>
OCR-related tasks	Text Captions (Sidorov et al., 2020)TextVQA (Singh et al., 2019)	22K	××	××	<i>v v</i>

Table 6: Description of datasets used in our alignment and VQA fine-tuning stages. Noting that in alignment process, we used 5K images from VQAv2 dataset.

	Dataset	Task	Split	Metric
Image-to-Text	MS-COCO (Lin et al., 2014) NoCaps (Agrawal et al., 2019) OK-VQA (Marino et al., 2019) GQA (Hudson and Manning, 2019)	Image captioning Image captioning VQA VQA	Test Val Val Test	CIDEr, BLEU, SPICE CIDEr, SPICE Accuracy Accuracy
Multimodal Generation	PhotoChat Zang et al., 2021	Image dialogue	Test	PPL, BLEU, ROUGE, FID

Table 7: Summary of the evaluation datasets and metrics.

# 971 eration and VQA tasks.

980

981

982

989

990

991

996

997

998

1001

InstructBlip (Dai et al., 2023) is a vision-language
instruction tuning framework that is trained with
BLIP-2 and capable of solving various visual language tasks.

976 MiniGPT-4 (Zhu et al., 2023) utilizes a single projection layer to align visual information from a pre978 trained vision encoder with an LLM. It employed
979 the same visual encoder as used in BLIP-2.

**LLaVA** (Liu et al., 2023b) employs a solitary projection layer to convert image features extracted from the pre-trained CLIP-ViT-L/14 visual encoder into the language embedding space of Vicuna.

Maria (Liang et al., 2021) is a neural conversation
agent which can leverage visual world experiences
sourced from a vast image index. It possesses the
ability to fetch a relevant image specific to the conversation and extract visual knowledge from it.

**Divter** (Sun et al., 2021) focuses on exploring multimodal dialogue generative models. Given the dialogue context, this model first generates a text response or image description and then generates an image according to the description.

## E Evaluation

For evaluating the quality of text generation, we utilize metrics such as BLEU, Rouge-L, Accuracy, and PPL (Perplexity). Additionally, following the approach of Vicuna (Chiang et al., 2023) and LLaVA (Liu et al., 2023b), we employ Chat-GPT to assess the generated responses from our model. Specifically, for the image captioning task,



Figure 6: Model's architecture for VQA finetuning. The module with blue background is referred to as BiDiffuser, while the rest is the architecture of MLLM using CLIP as the image encoder (such as LLaVA).

we randomly select 30 images from the MS-COCO Karpathy split and then let ChatGPT score the responses generated by EasyGen and the baseline models. ChatGPT evaluates the models' responses based on relevance, details, and accuracy and assigns an overall score between 1 and 10, with a higher score indicating better performance. To evaluate the quality of image generation, we use the Frechet Inception Distance (FID) score (Heusel et al., 2017), which measures the divergence between two multivariate normal distributions. 1002

1003

1004

1005

1006

1007

1008

1009

1010

1011

1012

1013

### **F** Implementation Details

**LLM** During the alignment process, we utilize 1014 the AdamW optimizer with  $\beta_0 = 0.9$ ,  $\beta_1 = 0.99$ , 1015

Task	Instruction Template
Image Captioning	USER: <image/> +random[query] Assistant:
LLaVA 80K	USER: Please answer question from this image: <image/> Question: <question> Assistant: USER: Image: <image/> Question: <question> Assistant: USER: Answer question <question> through the image <image/> Assistant:</question></question></question>
Multimodal Dialogue	USER: Dialog history+ <photo>+Dialogue history Assistant:</photo>
VQA	USER: Image: <image/> Question: <question> Short answer: Assistant: USER: Image: <image/> Question: <question> Answer the option's letter. Assistant:</question></question>

Table 8: Examples of task instruction templates. <image> represents the input image, <question> denotes the question in the VQA and LLaVA 80K dataset, and <photo> is the image description of the input image.

and weight decay of 0. The LLMs are trained with 1016 a cosine learning rate of 2e-5 and a warmup rate 1017 of 0.03. We use a batch size of 96 for the frozen 1018 1019 LLMs and 32 for the tuned LLMs. During training, we convert the LLMs (FlanT5XL/Vicuna-7B) to 1020 BFloat16/FP16 and BiDiffuser to FP16. During the 1021 VQA tuning process, we use CLIP ViT-L/14 336px as additional image encoder. We finetune EasyGen 1023 on mixture datasets for 1 epoch with a batch size 1024 1025 of 32. We adopt the AdamW optimizer with  $\beta =$ (0.9, 0.99) with the learning rate is 2e-5. We use 1026 1027 a cosine learning rate decay with a learning rate is 2e-5 and warmup ration is 0.03. 1028

> **Diffusion Module** We inherit the settings from UniDiffuser and utilize pre-trained weights from its checkpoint for our text-to-image generator. The model is fine-tuned on the MS-COCO and VG dataset, which contains images with a resolution of  $512 \times 512$ , for 10 epochs with a batch size of 312. For all of our sampling processes, we employ DPM-Solver with 50 steps.

1029

1030

1031

1033

1034

1035

1036

1037

1038

1039

1040

1041

1043

Provide a brief description of the given image.

Table 9: For the image captioning task, a query instruction is randomly selected.

## **G** Instruction Tuning

We list the instructions for different tasks in the main paper in Table 8. Specifically, the queries used to describe image contents are presented in Table 9. Table 8 shows the templates used in Vicuna, if the LLM is FlanT5, kindly use "Human" to substitute "USER" in the instruction templates. Model architecture for VQA finetuning is shown in1044Figure 6. EasyGen integrates the outputs of BiDif-<br/>fuser with images encoded by CLIP ViT-L/14. We1045freeze CLIP and BiDiffuser while only tuning the<br/>parameters of the LLM and projection layers.1047

1049

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1066

1067

1070

# H Training Efficiency

Table 10 summarizes the key factors in training EasyGen. The training process of EasyGen is computationally efficient, especially with the utilization of the parameter-efficient fine-tuning method LoRa (Hu et al., 2021). To enable multimodal response generation, we further train the aligned EasyGen. This process entails fine-tuning the LLM (FlanT5XL) on the PhotoChat dataset for 2 epochs, which typically requires approximately 4 A100 (80G) GPU hours.

## I Impact of Alignment Manners

In Table 11, we investigate the impact of different alignment manners on EasyGen. After removing the ITDM loss, the performance is slightly weaker than the original model. It is evident that the MSE Loss can help to align the semantic spaces of the two models. Furthermore, the performance of the model will drop significantly after removing the cross-entropy loss, suggesting that constraints via the language model play a key role.

## J More Qualitative Results

We present several instances on PhotoChat dataset 1071 in Figure 7 and the image-captioning task in Figure 8. In Figure 9, 11, 10, we compare EasyGen 1073 with state-of-the-art multimodal language models. 1074 The responses of MiniGPT-4, LLaVA, mPLUG-1075 owl and InstructBlip are obtained from their offi-1076 cial demos. Morever, in Figure 12, 13, we show 1077 EasyGen's ability to accept multimodal inputs and 1078 generate multimodal responses. 1079

<sup>1.</sup> Describe the image concisely.

<sup>3.</sup> Can you describe this image briefly?

Provide a summary of visual elements depicted in the image.
 Give me the essential characteristics of the photograph in a concise manner.

<sup>6.</sup> Rephrase the image depicted in a concise manner.

<sup>7.</sup> Describe the objects in this image no in detail.

<sup>8.</sup> Please introduce the image for me briefly.

<sup>9.</sup> Give me the image's short descriptions.

<sup>10.</sup> Please provide a general depiction of the image presented.

Model	Trainable Param.	Training Images	Training Cost
Pre-training BiDiffuser	952M	169K	120 (A100 80GB) GPU hours
Alignment			
Projection Layers + 🍀 T5XL	4M	163K	20 (RTX3090 24GB) GPU hours
Projection Layers + 😋 T5XL	3B	173K	20 (A100 80GB) GPU hours
Projection Layers + 💍 Vicuna 7B	7B	173K	72 (A100 80GB) GPU hours
Projection Layers + Ċ Vicuna 7B(LoRa)	610M	251K	20 (A100 80GB) GPU hours

Table 10: EasyGen's trainable parameters, training data size, and training cost during alignment process.

Model	NoCaps (val)		COCO (Karpathy)			OK-VQA	GQA
	CIDEr	SPICE	SPICE	BLEU@4	CIDEr	Accuracy	Accuracy
EasyGen Mid-Align FlanT5XL	121.2	15.5	25.1	43.5	145.7	31.5	22.6
+ w/o ITDM	118.6	15.3	24.8	42.2	141.5	-	-
+ w/o ITG	93.2	12.9	23.0	35.1	127.6	-	-

Table 11: Ablation studies on the instruction-tuning process and loss functions.



Figure 7: Examples of the generated responses on PhotoChat dataset. The text highlighted in red indicates the objects present in the image. The turns prefixed with A/B denote the given context.



EasyGen: A giraffe eating grass on a green grass field.



EasyGen: A snowboarder flying through the air on a frozen mountain slope.



EasyGen: A kitchen with a stove and oven in a fireplace.



EasyGen: Two cats separated by the window are looking forward at the same time.



EasyGen: Two children are riding on skis with their parents in a snowy street.



EasyGen: A long wooden bench with pictures of the world.

Figure 8: Examples of image captioning results by EasyGen.

#### < MultiModal Chat Group (((@)



Please tell me the information about this picture.

#### EasyGen (ours):

This painting is a classic representation of a young girl from the 17th century. The girl is depicted wearing a blue and white dress, with a pearl earring in her ear. The painting is a masterpiece of Dutch Golden Age art, a period of great cultural and economic prosperity in the Netherlands. The girl in the painting is likely depicted as a middleclass citizen, and the painting may have been commissioned as a portrait or for a specific purpose, such as a family crest or a religious painting. The painting is known for its attention to detail, vivid colors, and beautifully rendered textures. It is a true masterpiece of Dutch art, and a must- see for anyone interested in the history of art.

mPLUG-C

6

The image is a painting depicting a young girl, likely a Dutch girl in a yellow dress, with her hair tied up in braids and a gold earring. She is looking directly into the viewer's eyes, creating an intimate connection with the observer. The painting is set in an indoor scene with a table and chairs visible in the background. The girl is positioned in front of the table, and there is an additional person standing to her left. There are also several other people in various positions around the room, suggesting a social gathering or a family setting.

#### LLaVA

The image is a portrait of a woman wearing a blue headscarf, a yellow dress, and a pearl earring. She is looking to her left, and her outfit and accessories create a visually appealing contrast against the background. The woman's attire and appearance suggest a sense of style and elegance.

#### MiniGPT-4

The painting depicts a young woman with a pearl earring looking out into the distance. The woman is wearing a yellow dress with a blue and white headband. The background is a mix of blue and green, with a few trees visible in the distance. The overall mood of the painting is peaceful and contemplative. The painting is likely a representation of a woman reflecting on her thoughts or feelings.





Here are a few examples of paintings that are similar in style to the one you described: 1. "Girl with a Pearl Earring" by Johannes Vermeer; 2. "The Procuress" by Pieter de Hooch; 3. "The Artist's Mother" by Gerard van Honthorst .... These paintings are all from the 17th century and depict women or groups of people in various settings, often with a focus on everyday life. They are known for their attention to detail and realistic depiction of light and shadow.

LLoVA:
 I'm sorry, as an Al language model, I am not able to generate images. However, I can suggest some ideas for paintings that share a similar style to the portrait of the woman in the blue headscarf, yellow dress, and pearl earring by Johannes Vermeer:
 MiniGPT-4:
 Here are some pictures that resemble the painting, featuring young girls with their hair up in braids and wearing gold earrings: A young girl in a yellow dress and braided hair, looking directly into the viewer's eyes. Another young woman with her hair in similar braid styles, wears a white dress with gold accents, and has a gold necklace. Two girls in matching yellow outfits, with braiding and gold jewelry, standing together and

looking into each other's eyes, creating a warm and inviting atmosphere.

Figure 9: In this case study, for the first question, EasyGen can give an accurate answer including the background information of the image. With the image generation ability of BiDiffuser, EasyGen can generate visual responses.





Where can I see this object appeared in this image?

#### EasyGen (ours):

The Statue of Liberty is located in New York City, United States. It is a famous landmark and tourist attraction, and you can see it in person by visiting Liberty Island in New York Harbor.

#### mPLUG-Ow

You can see this blue and white 3D rendering of the Statue of Liberty located in the United States. The image features an artistically detailed representation of the famous landmark, which stands tall and proud in the skyline. The statue represents the ideals of freedom, democracy, and human rights. The scene is captured in a blue sky, emphasizing the beauty and grandeur of the iconic structure.

LLaVA

You can see the Statue of Liberty in the image, standing tall and proud in front of a blue sky.

#### I plan to travel here, what should I pay attention to?

#### EasyGen (ours)

When visiting the Statue of Liberty, there are a few things you should keep in mind: 1. The Statue of Liberty is a colossal neoclassical sculpture on Liberty Island in New York Harbor within New York City, United States. It was a gift from the people of France to the people of the United States and was dedicated on October 28, 1886, to commemorate the centennial of the signing of the United States Declaration of Independence. 2. The statue is 301 feet (92 meters) tall, including its pedestal and base, and is made of iron and steel. It was designed by the French sculptor Frederic Auguste Bartholdi and the American architect Emmanuel Louis Masqueray. 3. To get to the Statue of Liberty, you can take a ferry from Battery Park in Lower Manhattan. There are several ferry services that offer tours to the statue, and it's important to plan ahead and book your tickets in advance to avoid long lines. Once you're on Liberty Island, you can take an elevator or stairs to the top of the pedestal, where you'll have a stunning view of 4. New York Harbor and the surrounding city skyline. It's important to note that the pedestal is closed to visitors, but you can still get a close-up view of the statue from the observation deck. 5. Be sure to check the weather forecast before your visit, as the statue can be closed due to high winds or other inclement weather conditions. And don't forget to bring your camera to capture the stunning views and memorable moments! mPI UG-O As you're looking at a rendering of an iconic white and gold statue of a woman in a field, pay attention to the details of the sculpture, the sky, and the surroundings. The statue is situated in a green field, which provides a picturesque and serene atmosphere. The sky in the background is blue, suggesting a clear and sunny day. To fully appreciate the beauty and significance of the statue, it's crucial to observe its detail, size, and the overall composition of the image. The statue may depict a woman, an important figure in history, or a symbol of liberty and freedom. Additionally, pay attention to the field's surroundings, such as the grass, vegetation, and any other elements that contribute to the overall aesthetic. 0 When planning a trip to visit the Statue of Liberty, there are several things you should pay attention to: 1. Tickets: You will need to purchase tickets in advance to visit the Statue of Liberty and Ellis Island. You can do this through the official Statue of Liberty National Monument website or through authorized ticket sellers. 2. Accessibility: The Statue of Liberty is accessible via ferry from Battery Park in Lower Manhattan. There are also guided tours available, which can provide more information about the history and significance of the statue. 3. Closed areas: Some areas of the Statue of Liberty, such as the pedestal and the crown, are closed to visitors due to safety concerns. Make sure to check the official website for the most up-to-date information on accessibility and restrictions. 4. Clothing and footwear: Wear comfortable shoes and dress appropriately for the weather, as you will be walking and standing for extended periods of time

Figure 10: In this case study, for the first question, EasyGen can give an accurate answer, but the responses of the other two models are a bit biased. For the second question, EasyGen and LLaVA both give reasonable advice.

#### < MultiModal Chat Group 🧰



How can the girl improve her working environment?

#### EasyGen (ours)

(

6

0

In the image, the young girl is sitting in a room with a laptop in front of her, which suggests that she might be working or studying. To improve her working environment, she can consider several factors: 1. Ergonomics: Ensure that the laptop is placed at eye level or slightly below, and the desk is at a comfortable height for her to maintain a neutral posture. Adjustable furniture like a desk or a chair can help achieve a comfortable working height. 2. Lighting: Ensure that the room is well-lit, with natural light or artificial lighting to minimize eye strain and improve focus. 3. Organization: Keep the room clean and organized, with important items within easy reach. This can help maintain focus and prevent distractions. 4. Comfortable environment: Creating a comfortable environment, such as a dedicated workspace separate from her bedroom, can help establish boundaries and promote productivity. 5. Time management: Establish a routine and allocate specific time slots for studying or working on her laptop. This can help her stay on track and avoid procrastination. By implementing these suggestions, the girl can enhance her working environment and achieve better focus, productivity, and overall comfort while working on her laptop. To improve her working environment, the girl could consider several options: 1. Organize her space: Ensure the area where she is working is clean, uncluttered, and well-lit. Remove distractions such as toys or other items that might cause her to lose focus. Arrange her books and laptop in a manner that is easy to access and comfortable for her to work in. 2. Create a designated work area: Set up a designated space for her to work, such as a desk, table, or a specific corner in the living room. This helps her maintain her focus and minimize distractions. 3. Provide a comfortable chair and proper lighting: Select a chair with good back support and proper ergonomics to ensure she is comfortable while working. Lighting should be bright and evenly distributed to minimize shadows and glare. 4. Set reasonable work hours: Set a reasonable time limit for her to work, and let her know when it's time to take a break. This can help her maintain focus and minimize fatigue. By creating a well-organized, comfortable, and productive working environment, the girl can improve her focus and efficiency while minimizing distractions during her time spent studying and doing homework. MiniGP There are several ways the girl can improve her working environment. One way is to make sure her desk and chair are comfortable and adjustable to her height. Another way is to have a clean and organized workspace, so she can easily find what she needs. Additionally, she could have a good lighting system to help her see her work better. Finally, she could have a cushion or pillow to sit on to help reduce fatigue. The girl in the image can improve her working environment by organizing her surroundings and creating a designated study area. To achieve this, she could separate the laptop from other items in the room and create a consistent workspace. This can help her maintain focus and reduce distractions, which can ultimately enhance her productivity and creativity. Additionally, using a comfortable chair and positioning the laptop at a suitable height can help prevent eye strain and promote better posture during long

periods of work. Using a calm atmosphere with adequate lighting and placing the laptop away from windows or other sources of glare can also aid in ensuring a more conducive working environment.

Figure 11: From this example, we can find that the response from EasyGen is more comprehensive and coherent. This shows EasyGen can give reasonable suggestions based on the given image.



Figure 13: Example of multimodal response generation.