Pre-trained Language Models Do Not Help Auto-regressive Text-to-Image Generation

Anonymous ACL submission

Abstract

Recent advances in image tokenizers, such as VQ-VAE, have enabled text-to-image generation using auto-regressive methods, similar to language modeling. However, these methods have yet to leverage pre-trained language models, despite their adaptability to various downstream tasks. In this work, we explore this gap by adapting a pre-trained language model for auto-regressive text-to-image generation, and find that pre-trained language models offer limited help. We provide a two-fold explanation by analyzing tokens from each modality. First, we demonstrate that image tokens possess significantly different semantics compared to text tokens, rendering pre-trained language models no more effective in modeling them than randomly initialized ones. Second, the text tokens in the image-text datasets are too simple compared to normal language model pre-training data, which causes the catastrophic degradation of language models' capability.

1 Introduction

Recent works in text-to-image generation primarily employ two kinds of methods: diffusion models (Ramesh et al., 2022; Saharia et al., 2022; Rombach et al., 2022) and auto-regressive models (Ramesh et al., 2021; Yu et al., 2022b). The latter is facilitated by "image tokenizers", such as VQ-VAE (van den Oord et al., 2017; Razavi et al., 2019) and VQ-GAN (Esser et al., 2021; Yu et al., 2022a), which transform an image into a sequence of discrete tokens, similar to text tokens (Figure 1 Left). Consequently, image and text tokens can be jointly modeled using auto-regressive algorithms like the Transformer (Vaswani et al., 2017) (Figure 1 Right).

The superiority of diffusion-based models when compared with auto-regressive-based methods for text-to-image generation still remains unclear. OpenAI's pioneering work, DALL-E (Ramesh et al., 2021), showcased the potential of auto-regressive methods in this domain. Yet, its successor, DALL-E 2 (Ramesh et al., 2022), transitioned to a diffusion-based architecture and achieved enhanced image generation quality. Later, Google released Imagen (Saharia et al., 2022) (diffusion-based) and Parti (Yu et al., 2022b) (auto-regressive-based) at the same time and demonstrated their comparable generation quality. Similarly, the retrieval-augmented methods, Re-Imagen (Chen et al., 2022) (diffusion-based) and CM3leon (Yu et al., 2023b) (auto-regressive-based), display similar performance in text-to-image generation tasks. A comparison based on zero-shot FID (Heusel et al., 2017) on the COCO dataset (Lin et al., 2014) can be found in Figure 2.

While these two approaches achieve similar performance, it is intriguing that *diffusion-based models consistently utilize pre-trained text encoders, whereas their auto-regressive counterparts generally do not*. For instance, Imagen (Saharia et al., 2022) (diffusion-based) reports that employing a stronger pre-trained text encoder, specifically T5 (Raffel et al., 2020), yields substantial improvements to using CLIP (Radford et al., 2021). Furthermore, they observe that scaling up the T5 text encoder leads to more pronounced improvements than scaling up the diffusion models. Conversely, Parti (Yu et al., 2022b) (auto-regressive-based) shows that using a pre-trained text encoder does not necessarily improve image quality in its Appendix. However, Parti employs an encoder-decoder architecture and uses BERT (Devlin et al., 2019), a relatively inferior text encoder, to initialize the encoder only. It remains unclear whether a decoder-only approach would benefit from recent advances in large language models (LLMs), given the clear similarity between language modeling and auto-regressive text-to-image generation.

In this work, we explore the potential of pre-trained LLMs for auto-regressive text-to-image generation. To enable the model to process both text and image tokens, we expand the size of the

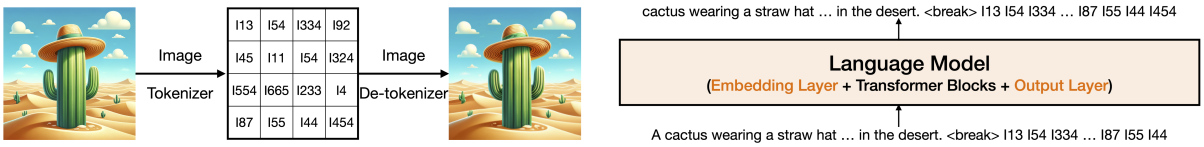

Figure 1: *Adapting language models for auto-regressive text-to-image generation. (Left)* An image is fed into an image tokenizer (MoVQGAN (Zheng et al., 2022)) and converted to a grid of discrete tokens, and it can be well-reconstructed with these image tokens. *(Right)* As images are converted to tokens similar to text tokens, we can enable language models to generate images by adapting its embedding layer and output layer.

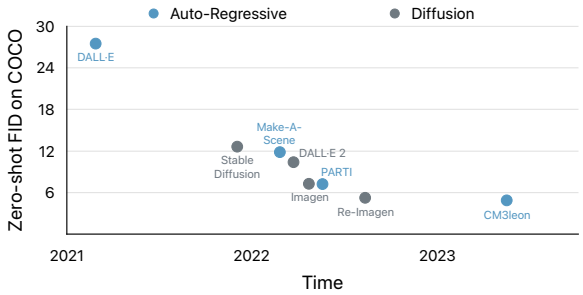

Figure 2: *Auto-regressive and diffusion based models achieve similar performances on text-to-image generation.* However, while all the diffusion models leverage pre-trained language models, all the auto-regressive models do not.

embedding and output layers by incorporating an image vocabulary from the image tokenizer. We initialize these added weights either randomly or using a novel contrastive alignment (elaborated later in Section 3.2), while the remaining weights are directly copied from the original models. Subsequently, we fine-tune the model on image-caption datasets, as depicted in Figure 1 Right.

Surprisingly, the results show that pre-trained language models achieve the same loss and image generation quality as the model that is entirely randomly initialized and trained from scratch (Figure 3). Furthermore, we observe a catastrophic deterioration in the model's text capabilities, such as world knowledge or in-context learning, after only minimal steps of fine-tuning (Table 1).

To understand this phenomenon, we break down the cross-entropy loss on image and text tokens, and find that 1) the loss on image tokens is the same between the pre-trained and randomly initialized model, and 2) the loss on text tokens of the pre-trained model is significantly lower at the beginning compared to the randomly initialized models, but the gap soon disappears after training (Figure 4).

The first finding of the loss on the image tokens is particularly interesting. We hypothesize that image tokens obtained from image tokenizers might either lack semantics or possess significantly different semantics compared to text tokens, which renders language pre-training not transferable to the image modeling task. To verify this hypothesis, we conduct unconditional image generation experiments by training the model on image tokens only. Our results show that 1) the pre-trained model achieves the same loss as the randomly initialized model, and 2) freezing any part of the pre-trained model results in a loss degradation (Figure 5). These indicate that optimal weights for language and image modeling are fundamentally different, making language pre-training not transferable to image modeling.

In summary, we share our experimental findings about pre-trained language models do not help auto-regressive text-to-image generation, and offer an explanation: 1) the intrinsic differences between image and text tokens make language pre-training ineffective for the image token modeling, and 2) the disproportionate ratio between image and text tokens (usually 30:1 for image-caption datasets) minimizes the impact of loss on text tokens and leads to catastrophic forgetting.

2 Pre-trained Language Models Do Not Help Text-to-Image Generation

2.1 Experimental Setup

Language model. We use the publicly available `open_lm` codebase and its `open_lm-1b` model for our experiments (Gururangan et al., 2023). This language model contains ∼1B parameters and is trained on 1.6T tokens on a mix of RedPajama (Computer, 2023), Pile (Gao et al., 2020), S2ORC (Lo et al., 2020), The Pile of Law (Henderson et al., 2022), Deepmind Math (Saxton et al., 2019), and RealNews (Zellers et al., 2019b). It achieves better or comparable performance compared to models with similar size such as OPT-1.3B (Zhang et al., 2022), Pythia-1B (Biderman et al., 2023), Neox-1.3B (Black et al., 2022), OPT-IML-1.3B (Iyer et al., 2022) on an average of 11 tasks such as HellaSwag (Zellers et al., 2019a) and MMLU (Hendrycks et al., 2021). More details can be found in the `open_lm` repository (Gururangan et al., 2023).

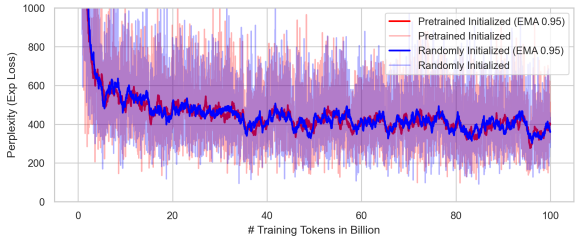

Figure 3: *Pre-trained language models do not help auto-regressive text-to-image generation.* Models are trained on the HQITP-134M image-caption dataset with 64 A100 80GB GPUs using batch size 1M tokens. EMA is Exponential Moving Average.

Image tokenizer. We use SBER-MoVQGAN (Zheng et al., 2022) as the image tokenizer, which is the current state-of-the-art publicly available image tokenizer that achieves 0.686 FID on Imagenet image reconstruction. Given an image with 256 $\times$ 256 resolution, it converts an image to 1,024 tokens with a vocabulary size of 16,384. Figure 1 (Left) shows a real reconstruction example from this tokenizer.

Dataset. For multi-modal training, we use an internal dataset referred to as High Quality Image-Text Pairs (HQITP) (Ranasinghe et al., 2023), which contains 134M high-quality image-caption pairs. We pre-process the dataset before training. Each image is center-cropped to 256 $\times$ 256 and converted to 1,024 tokens. Each caption is tokenized with NeoX tokenizer with an average of 30 tokens. We add six special tokens corresponding to the beginning and end of document, text segment, and image, respectively. This results in input sequences of the form "<doc> <text> ...text tokens... </text> <image> ...image tokens... </image> </doc>", and pad them into 1,152 tokens with the special <pad> token.

Training setups. Models are trained with 100B tokens using 64 A100 80GB GPUs with batch size 1M tokens. We use the AdamW (Loshchilov and Hutter, 2019) optimizer with a cosine learning rate schedule with 2K warm-up steps and a peak learning rate of 0.0003. This mimics the settings reported in (Aghajanyan et al., 2023). We also tried different hyper-parameters, such as learning rates from 0.00005 to 0.0003 and batch size from 0.5M to 2M tokens, and found no significant influences on the conclusions.

2.2 Results

In Figure 3, we present the perplexity (exponential of loss) during training for both the pre-trained and randomly initialized models. Intriguingly, across the entire 100B token training regimen, the loss of the pre-trained model aligns closely with that of the randomly initialized one. Beyond this, a sharp decline in text capabilities of the pre-trained model is observed after training on 5B tokens, as illustrated in Table 1. At this point, both the model's world knowledge and its in-context learning ability are entirely diminished.

Original Completion	Completion after Training 5B Tokens
Simply put, the theory of relativity states that **the speed of light is the same for all observers, regardless of their location in the universe.**	Simply put, the theory of relativity states that **iles must be able to see the invisible.**
Translate English to French: sea otter => loutre de mer peppermint => menthe poivrée plush girafe => girafe peluche cheese => **fromage**	Translate English to French: sea otter => loutre de mer peppermint => menthe poivrée plush girafe => girafe peluche cheese => **I love cheese**

Table 1: *Concrete examples of forgetting.* We observe a severe deterioration of the model's language capability, such as knowledge and in-context learning, after a small amount of training. Model completions are bolded.

To delve deeper into this phenomenon, we separate the cross-entropy loss into two components: text tokens and image tokens, displayed separately in Figure 4. As anticipated, the pre-trained model begins with a significantly lower text loss in comparison to its randomly initialized counterpart. Yet, due to the overwhelming image-text token ratio (30:1), this initial advantage is obscured in the aggregate loss. Furthermore, any benefit the pre-trained model offers in text loss diminishes soon during training. In contrast, for image tokens, there is no difference between the pre-trained and randomly initialized models. We hypothesize that the inability of effectively transferring a pre-trained language model to image token modeling is caused by the distinction between image and text tokens.

Moreover, loss on text tokens is substantially lower than image tokens, and even lower than typical language models trained on text-only data. This is because texts in image-caption datasets such as HQITP are less complex than those in standard text-only pre-training corpora, which also explains the catastrophic degradation of the model's text capability.

3 Image Tokens Are Drastically Different From Text Tokens

Why there is no difference between the loss of pre-trained and randomly initialized models on the

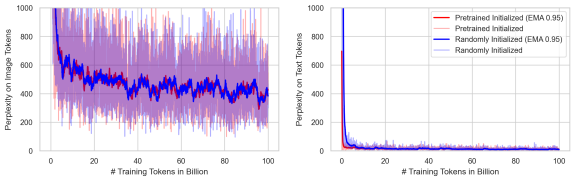

Figure 4: *Break-down loss on image and text tokens.* Models are trained on the HQITP-134M image-caption dataset with 64 A100 80GB GPUs using batch size 1M tokens.

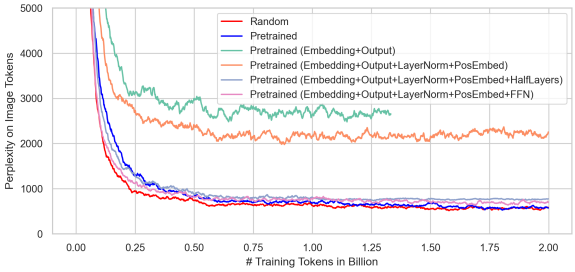

Figure 5: *Pre-trained language models do not help to model image tokens.* Models are trained only on the HQITP dataset's image tokens without any text tokens. We also compare the full fine-tuning with electively fine-tuning components of the pre-trained models (shown in parenthesis). EMA 0.95 is applied to the plot.

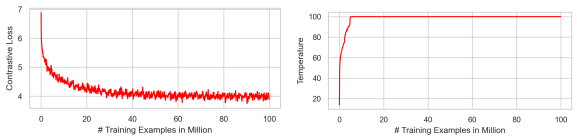

Figure 6: *Image-text token contrastive alignment. (Left)* The contrastive loss plateaus quickly, indicating a difficulty in aligning text and image tokens directly at a bag-of-words level. *(Right)* The learnable temperature in the contrastive loss during training for reference.

image tokens? We hypothesize image tokens are significantly different from text tokens, for example, they lack semantics or have drastically different semantics compared to text tokens, which makes the pre-trained language model not transferable to image token modeling. Our unconditional image generation and image-token alignment experiments verify this hypothesis.

3.1 Unconditional Image Generation

To assess if pre-trained language models benefit image tokens, we perform unconditional image generation experiments. Unlike the text-to-image generation setup, we removed all text tokens, leaving only the image tokens. This approach rigorously examines if image tokens benefit from pre-trained language models. As shown in Figure 5, pre-trained language models yield the same loss as models initialized randomly.

Additionally, we selectively tune components of the pre-trained models: 1) only the embedding and output layer; 2) 1 plus layer norm and positional embedding; and 3) 2 plus the first half of layers; 4) 2 plus the feed-forward layers (FFN). Figure 5 presents these loss metrics. The findings reveal that none of these configurations achieves as low a loss as a fully tunable model. This underscores the divergence in optimal weights for modeling text and image tokens, suggesting that any part of the text-trained weights is sub-optimal to transfer to image tokens.

3.2 Image-Text Token Contrastive Alignment

To understand whether image tokens have similar semantics as text tokens, we aligned image tokens with text tokens using a contrastive approach, inspired by methods like CLIP (Radford et al., 2021). Given an image, we tokenize it into 1024 tokens and compute its bag-of-words image embeddings as its representation. Similarly, we tokenize the corresponding caption and compute its bag-of-words text embeddings. The text embeddings are initialized from a pre-trained language model while the image embeddings are randomly initialized. For a batch of $N = 1024$ image-caption pairs, the contrastive objective from CLIP is employed to maximize the cosine similarity between matched image-caption l_2-normalized representations and to minimize the similarity for non-matching pairs. Only the image embeddings are updated during training.

In Figure 6, we illustrate that the contrastive loss plateaus quickly, indicating a difficulty in aligning text and image tokens directly at a bag-of-words level. Indeed, after training, when querying the closest text tokens for any image token, we observe that they predominantly align with noisy, semantically void text tokens. Furthermore, when we use the trained image embeddings as initialization for text-to-image generation, as opposed to random initialization, there is no discernible improvement.

4 Conclusion

This study highlights the difficulty of naively adapting a text-only language model to handle multi-modal contents, such as texts and images. Given the challenge of the disparities between image tokens and text tokens, a valuable avenue for future experiments is to employ tokenizers that align semantically with text tokens, such as SEED (Ge et al., 2023) or SPAE (Yu et al., 2023a).

Limitations

Our study has some limitations. First, the results are based on the VQGAN image tokenizer, which does not align semantics between image tokens and text tokens. Tokenizers that semantically align image tokens with text tokens might yield different outcomes. Second, we observed severe degradation in language model capabilities during fine-tuning, suggesting that exploring methods to avoid catastrophic forgetting could be a promising future research direction. Additionally, the experiments required extensive computational resources, which might limit reproducibility. Despite these limitations, our findings remain useful and provide valuable information for future research.

References

Armen Aghajanyan, Lili Yu, Alexis Conneau, Wei-Ning Hsu, Karen Hambardzumyan, Susan Zhang, Stephen Roller, Naman Goyal, Omer Levy, and Luke Zettlemoyer. 2023. Scaling laws for generative mixed-modal language models. In *ICML*.

Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O'Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, et al. 2023. Pythia: A suite for analyzing large language models across training and scaling. In *ICML*.

Sidney Black, Stella Biderman, Eric Hallahan, Quentin Anthony, Leo Gao, Laurence Golding, Horace He, Connor Leahy, Kyle McDonell, Jason Phang, Michael Pieler, Usvsn Sai Prashanth, Shivanshu Purohit, Laria Reynolds, Jonathan Tow, Ben Wang, and Samuel Weinbach. 2022. GPT-NeoX-20B: An open-source autoregressive language model. In *ACL Workshop*.

Wenhu Chen, Hexiang Hu, Chitwan Saharia, and William W Cohen. 2022. Re-imagen: Retrieval-augmented text-to-image generator. In *ICLR*.

Together Computer. 2023. Redpajama: An open source recipe to reproduce llama training dataset.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*.

Patrick Esser, Robin Rombach, and Björn Ommer. 2021. Taming transformers for high-resolution image synthesis. In *CVPR*.

Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. 2020. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*.

Yuying Ge, Yixiao Ge, Ziyun Zeng, Xintao Wang, and Ying Shan. 2023. Planting a seed of vision in large language model. *arXiv preprint arXiv:2307.08041*.

Suchin Gururangan, Mitchell Wortsman, Samir Yitzhak Gadre, Achal Dave, Maciej Kilian, Weijia Shi, Jean Mercat, Georgios Smyrnis, Gabriel Ilharco, Matt Jordan, Reinhard Heckel, Alex Dimakis, Ali Farhadi, Vaishaal Shankar, and Ludwig Schmidt. 2023. open_lm: a minimal but performative language modeling (lm) repository. GitHub repository.

Peter Henderson, Mark Krass, Lucia Zheng, Neel Guha, Christopher D Manning, Dan Jurafsky, and Daniel Ho. 2022. Pile of law: Learning responsible data filtering from the law and a 256gb open-source legal dataset. In *NeurIPS*.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. In *ICLR*.

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. In *NeurIPS*.

Srinivasan Iyer, Xi Victoria Lin, Ramakanth Pasunuru, Todor Mihaylov, Daniel Simig, Ping Yu, Kurt Shuster, Tianlu Wang, Qing Liu, Punit Singh Koura, et al. 2022. Opt-iml: Scaling language model instruction meta learning through the lens of generalization. *arXiv preprint arXiv:2212.12017*.

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 *ECCV*.

Kyle Lo, Lucy Lu Wang, Mark Neumann, Rodney Kinney, and Daniel Weld. 2020. S2ORC: The semantic scholar open research corpus. In *ACL*.

Ilya Loshchilov and Frank Hutter. 2019. Decoupled weight decay regularization. In *ICLR*.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning transferable visual models from natural language supervision. In *ICML*.

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *JMLR*.

Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*.

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 *ICML*.

Kanchana Ranasinghe, Brandon McKinzie, Sachin Ravi, Yinfei Yang, Alexander Toshev, and Jonathon Shlens. 2023. Perceptual grouping in contrastive vision-language models. In *ICCV*.

Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. 2019. Generating diverse high-fidelity images with VQ-VAE-2. In *NeurIPS*.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *CVPR*.

Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Raphael Gontijo-Lopes, Burcu Karagol Ayan, Tim Salimans, et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. In *NeurIPS*.

David Saxton, Edward Grefenstette, Felix Hill, and Pushmeet Kohli. 2019. Analysing mathematical reasoning abilities of neural models. In *ICLR*.

Aäron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. 2017. Neural discrete representation learning. In *NeurIPS*.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NeurIPS*.

Jiahui Yu, Xin Li, Jing Yu Koh, Han Zhang, Ruoming Pang, James Qin, Alexander Ku, Yuanzhong Xu, Jason Baldridge, and Yonghui Wu. 2022a. Vector-quantized image modeling with improved VQGAN. In *ICLR*.

Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, et al. 2022b. Scaling autoregressive models for content-rich text-to-image generation. *TMLR*.

Lijun Yu, Yong Cheng, Zhiruo Wang, Vivek Kumar, Wolfgang Macherey, Yanping Huang, David A Ross, Irfan Essa, Yonatan Bisk, Ming-Hsuan Yang, et al. 2023a. Spae: Semantic pyramid autoencoder for multimodal generation with frozen llms. *arXiv preprint arXiv:2306.17842*.

Lili Yu, Bowen Shi, Ramakanth Pasunuru, Benjamin Muller, Olga Golovneva, Tianlu Wang, Arun Babu, Binh Tang, Brian Karrer, Shelly Sheynin, et al. 2023b. Scaling autoregressive multi-modal models: Pretraining and instruction tuning. *arXiv preprint arXiv:2309.02591*.

Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019a. HellaSwag: Can a machine really finish your sentence? In *ACL*.

Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. 2019b. Defending against neural fake news. In *NeurIPS*.

Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer, Kurt Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettlemoyer. 2022. Opt: Open pre-trained transformer language models. *Preprint*, arXiv:2205.01068.

Chuanxia Zheng, Tung-Long Vuong, Jianfei Cai, and Dinh Phung. 2022. Movq: Modulating quantized vectors for high-fidelity image generation. In *NeurIPS*.