

Predicting Winning Captions for Weekly New Yorker Comics

Stanley Cao
Stanford University
Department of Computer Science
stancao@stanford.edu

Sonny Young
Stanford University
Department of Chemistry
ssyoung@stanford.edu

Abstract

Image captioning using Vision Transformers (ViTs) represents a pivotal convergence of computer vision and natural language processing, offering the potential to enhance user experiences, improve accessibility, and provide textual representations of visual data. This paper explores the application of image captioning techniques to New Yorker cartoons, aiming to generate captions that emulate the wit and humor of winning entries in the New Yorker Cartoon Caption Contest. This task necessitates sophisticated visual and linguistic processing, along with an understanding of cultural nuances and humor. We propose several new baselines for using vision transformer encoder-decoder models to generate captions for the New Yorker cartoon caption contest.

1. Introduction

Image captioning using Vision Transformers (ViTs) is a pivotal area of research, merging the domains of computer vision and natural language processing. The ability to automatically generate accurate and contextually relevant descriptions for images comes with numerous benefits, including enhancing user experience in various digital platforms, increasing accessibility for visually impaired users, and providing textual representations of visual data.

In this paper, we explore the complexities of applying image captioning techniques to New Yorker cartoons, particularly focusing on generating captions that closely emulate the wit and humor of winning entries in the New Yorker Cartoon Caption Contest. This task requires not only a sophisticated combination of visual and linguistic processing, but also the ability to caption images in a completely orthogonal way: the winning caption rarely describes the cartoon, but rather *presupposes* a deep understanding of the scene along with the concomitant cultural references. By aiming to create models that can effectively mimic human-like humor in captioning, we

demonstrate potential advancements in the development of AI systems that have more holistic multi-modal understanding.

2. Related Work

Image captioning has so far been applied to numerous generative and classification problems dealing with satellite data [22], pathological diagnosis [23, 17], or real-life object detection [1, 28, 10] that primarily deal with natural images. Much success in these tasks has been attributed to attention models popularized by transformers [24] which embed sequential data in a parallel and expressive manner, and more recently, ViTs which are considered the state of the art architecture for image recognition [4].

The challenge is especially apparent when dealing with artificially constructed images, such as cartoons or comics. These images often contain abstract, stylized representations that may not correspond directly to real-world objects and may embed complex narratives and humor, which are particularly difficult for AI to parse and understand. Furthermore, ViTs are known to require significant pre-training data, close to 100 million images before its performance is comparable to a CNN [4]. Some have attempted to circumvent the large dataset issue by training ViTs without using natural data [14]. One particularly noteworthy model is LLaVA, a vision language model (VLM) that is capable of incorporating multi-modal information in its generated output [13]. [11] showcases that with simple modifications to the multi-modal projection layer and further finetuning on VQA datasets, the model is able to achieve improved performance on various downstream tasks. However, evaluating the model beyond traditional question-answering tasks remains largely unexplored.

Natural Language Inference (NLI) is a key area that can enhance a model’s ability to interpret language, which is crucial for improving how AI systems grasp concepts like humor [2], common sense [25], and overall sentence interpretation [29]. These developments sug-



Image Description	Uncanny Description	Entities	Caption
Three dogs are in a field in front of some rolling hills. One of the dogs is holding a stick and winding up, as if it's about to hit a baseball, while the other two dogs look on.	It's unusual that the dog is playing a sport like baseball because dogs don't have the cognitive capacity to play that human game.	Fetch, Dog, Anthropomorphism.	He's his own best friend.

Figure 1: Detailed Information including selected metadata entries for New Yorker Cartoon Caption Contest 102

gest a path forward for refining AI approaches to better understand and engage with human communication nuances.

We focus on the New Yorker Cartoon Caption Contest, which requires humorous captioning that transcends pure information content. Each week, participants are invited to submit captions for a cartoon, often characterized by its humorous, satirical, or ironic undertone, making the task of generating a fitting caption not just a test of understanding visual elements, but also of grasping subtle cultural and contextual nuances. Successfully modeling the New Yorker Cartoon captions would demonstrate a significant leap towards AI systems that can understand and generate human-like text responses.

Notably, [5] has explored this area of research, not only creating a dataset of New Yorker cartoons from prior contests, but proposing three different tasks to benchmark the capabilities revolving around multi-modal humor understanding: matching, quality ranking, and explanation. The authors report GPT-4 having superior performance in matching and quality ranking, achieving 84.5% accuracy in matching a true caption among a pool of 5 fake captions. GPT-4 significantly outperformed all other methods, but it was unable to explain jokes as well as humans. The tasks proposed in [5] are largely classification tasks, with the exception of their proposed explanation task. Since LLMs are unable to achieve human-level performance, these results suggest that there is room for further exploration in the generation of humorous captions.

3. Dataset

We utilized the New Yorker Caption Contest dataset available on Hugging Face [7, 18, 21]. This dataset consists of cartoons that are used for three different tasks to assess the multi-modal capabilities of ViTs and encoder-decoder models. Since the task is to generate captions from the cartoon alone, much like what is expected of humans for the caption contest, we make use of the dataset that corresponds to the explanation task proposed in [5], which not only contains an cartoon and winning caption pair, but also important metadata such

as human-annotated descriptions of the cartoon, an explanation for why the caption is humorous, important entities that are relevant to the cartoon, and leading questions that appear to motivate the winning caption. Figure 1 shows a single data entry with selected metadata.

This dataset comprises approximately 2.6K data entries. We use the same train-validation-test splits as provided in [7], with about 2.34k entries in the train set and about 130 entries in the validation and test set. We do not preprocess or normalize the cartoons in advance, as the ViTs used in this paper contain their own custom image transformations / processing techniques (see Section 4 for more details). As all cartoons are of different sizes, we have varying image resolutions, though most of ViTs presented in this paper normalize image resolution by padding the patches derived from the partitioning the image.

4. Methodology and Experiments

As our overarching goal is to generate the caption from the cartoon alone, we employ multi-modal vision encoder-decoder models for this task.

As one of the main findings in [5], computer vision serves as a bottleneck for top-quality explanation generation; models that have access to human-labeled cartoon scene descriptions perform significantly better. However, providing the model with both cartoon and human-annotated metadata does not accurately mimic the task of the New Yorker Cartoon Caption Contest — the metadata for a given cartoon (e.g. image description) is inherently up for interpretation. It is the creative task of determining which aspects of the cartoon to satirically emphasize that is highlighted in such a contest. Thus, providing the model with metadata, which likely improves performance, means that we are leaking information to the model that a normal human participant in the contest does not have. In the remainder of this section, we outline 8 different training pipelines we devised for this task. We used the Hugging Face and Pytorch libraries for our tasks [27, 16]

4.1. CLIP + GPT2

For our baseline, we implemented a CLIP-GPT2 encoder-decoder model that takes a cartoon as input and is required to output the caption. We employ the CLIP vision transformer model [19] to map all input cartoons into a multi-modal embedding space. From this embedding space, we use GPT2 [20] as a decoder, translating this image embedding into a cartoon caption by using cross-attention between the encoder’s output and the decoder’s caption. In the experiments described in this section, we finetune our model end-to-end on the dataset for 10 epochs, using the AdamW optimizer with a learning rate of 5e-5.

4.1.1 Caption Only

In this formulation, we ignore all metadata, simply tasking the model to generate the caption from the image alone. We include this approach as a baseline for this task.

4.1.2 Caption + Metadata (all)

In this formulation, we prepend the metadata to the target label, so that the caption is generated last. This training setup encourages the model to learn to describe the image and generate relevant metadata first. Due to the transformer architecture in decoder models, the model would use its attention mechanism to refer to previously self-generated metadata, culminating in a generated caption that is well-informed of the prior metadata. Figure 2 shows an example of synthesized target label.

Label: scene: a field description:
3 dogs are out in a field.
One is standing in front of
the other two and preparing
to throw a stick uncanny:
It’s unusual to see dogs
holding items as they do
not have oposable digits
entities: Anthropolomorphism,
Dog, Baseball. caption:
He’s his own best friend.

Figure 2: Modified target label with metadata prepended for the cartoon in Figure 1

4.2. LLaVA-NeXT

LLaVA-based models are variants of vision encoder-decoder models, which projects the input image onto an embedding space that has the same dimension as the language embedding space. By training this light-weight

projection layer, LLaVA-NeXT is able to turn images into tokens, which are then subject to the same attention mechanisms as the textual data.

In [13], they used a *CLIP-ViT-L/14* to *Vicuna* LLM encoder-decoder model. However, in our experiments, we use the LLaVA-NeXT model. In the original paper, the authors use a *CLIP-ViT-L-336px* to *Vicuna* LLM [12], but we make use of a variant hosted on Hugging Face, which replaces *Vicuna* with *Mistral* LLM.

4.2.1 0-shot Setting

We evaluate the efficacy of LLaVA-NeXT on this task in a 0-shot setting. Figure 3 shows the 0-shot prompt we used to query the model, and we extract its response as the generated caption. The model is not presented any examples in advance, and must execute the task using only its prior knowledge base.

Input: [INST] <image> Please write
a winning caption for the
provided cartoon in the
New Yorker caption contest.
[/INST]

Figure 3: 0-shot prompt for LLaVA-NeXT and GPT-4V. The <image> token refers to the placement of the image embedding for the LLaVA-NeXT Model.

4.2.2 5-shot Setting

In addition to the 0-shot setting, we evaluate the efficacy of LLaVA-NeXT on this task in a 5-shot setting. Specifically, we show the model 5 examples, as demonstrated in Figure 4, with the intention to encourage the model to learn the style of the task.

Input: [INST] <image1> [/INST] CAPTION 1
Input: [INST] <image2> [/INST] CAPTION 2
Input: [INST] <image3> [/INST] CAPTION 3
Input: [INST] <image4> [/INST] CAPTION 4
Input: [INST] <image5> [/INST] CAPTION 5
Input: [INST] <image> Please write a
winning caption for the provided
cartoon in the New Yorker caption
contest. [/INST]

Figure 4: 5-shot prompt for LLaVA-NeXT. The <image> token refers to the placement of the image embeddings. The final user input asks the model to generate a caption after it has seen 5 previous examples of human-written winning captions.

4.2.3 Chain-of-Thought Prompting

Chain-of-Thought (CoT) is a prompt engineering technique designed to enhance language models' performance on logic, calculation, and decision-making tasks by structuring the input prompt to simulate human reasoning [26]. We provide LLaVA-NeXT with a series of intermediate reasoning steps that encourages more sophisticated decision-making. Figure 5 showcases our proposed Chain-of-Thought prompting technique.

4.2.4 Finetuning with QLoRA

We finetuned LLaVA-NeXT on our training set in a question-answer format to remain consistent with the model's prior pre-training. Low-Rank Adaptation (LoRA) is a popular finetuning mechanism for LLMs that significantly reduces the number of trainable parameters without losing the benefit of downstream finetuning [6]. LoRA approximates updates to weight matrices with a low-rank matrix decomposition. Thus, for a given weight matrix $W_0 \in \mathbb{R}^{d \times k}$, LoRA constrains the update ΔW to be approximated by a low-rank matrix decomposition $\Delta W \approx BA$, where $B \in \mathbb{R}^{d \times r}$ and $A \in \mathbb{R}^{r \times k}$, and $r \ll \min(d, k)$. This means that the update to weight matrices becomes

$$h = (W_0 + \Delta W)x \approx (W_0 + BA)x.$$

During training, W_0 is frozen and does not receive gradient updates, while A and B are trainable parameters.

Due to compute limitations, we use the 7B parameter model as opposed to the 34B parameter model. In addition, we implemented quantized LoRA (QLoRA) and LoftQ initialization [3, 8] to further optimize memory usage. We finetuned the LLaVA-NeXT model end-to-end on an A100 80GB GPU for 10 epochs with an r value of 32, learning rate of $1e-4$, α of 16 on all linear layers.

4.3. GPT-4V

We repeated the same evaluation of this task using the GPT-4V multi-modal model with both 5-shot and complex CoT mechanisms. We used the OpenAI API to query GPT-4V and utilized the same prompts as our evaluation for LLaVA-NeXT.

5. Results

We compare our generated captions to the human-written captions via automated evaluation metrics (e.g., BLEU and ROUGE) and manual quality examination.

5.1. BLEU and ROUGE Scores

A well-known metric for automated quality evaluation of machine generated text is the Bilingual Eval-

Input: [INST] <image> Please write an engaging, witty entry for the New Yorker Caption Contest based on this image. You are competing against 5,000 entrants from across the world, many of whom have made hundreds of previous submissions.

Please follow these rules:

1. First, describe the comic in as much objective detail as you can, subject to a limit of 4 sentences max. Leave no stone unturned and describe all main objects and characters.

2. Then, list as many uncanny aspects of the comic as you can think of (list at least 5). Then, for each one, explain why it is uncanny and why the comic artist likely inserted it. Think about what makes each aspect of the unusual occurrence humorous.

3. Based on your descriptions above and using great taste, come up with 5 potential caption concepts, each of which should be your best work and maximally diverse. Then, roast each one in 2-3 sentences.

4. Write 5 improved concepts after that, explaining why each is better.
[/INST]

LLM: GENERATED RESPONSE

Input: [INST] Thanks! Now, select the best caption from all the versions you've written above. Write only the caption as a single line, with no other text.
[/INST]

Figure 5: Chain-of-Thought prompting for both LLaVA-NeXT and GPT-4V.

uation Understudy (BLEU) score metric, which is designed to evaluate the quality of text from one natural language to another[15, 9]. All models performed poorly with respect to the BLEU score metric. Recall-Oriented Understudy for Gisting Evaluation (ROUGE) is another popular scoring metric for evaluating auto-

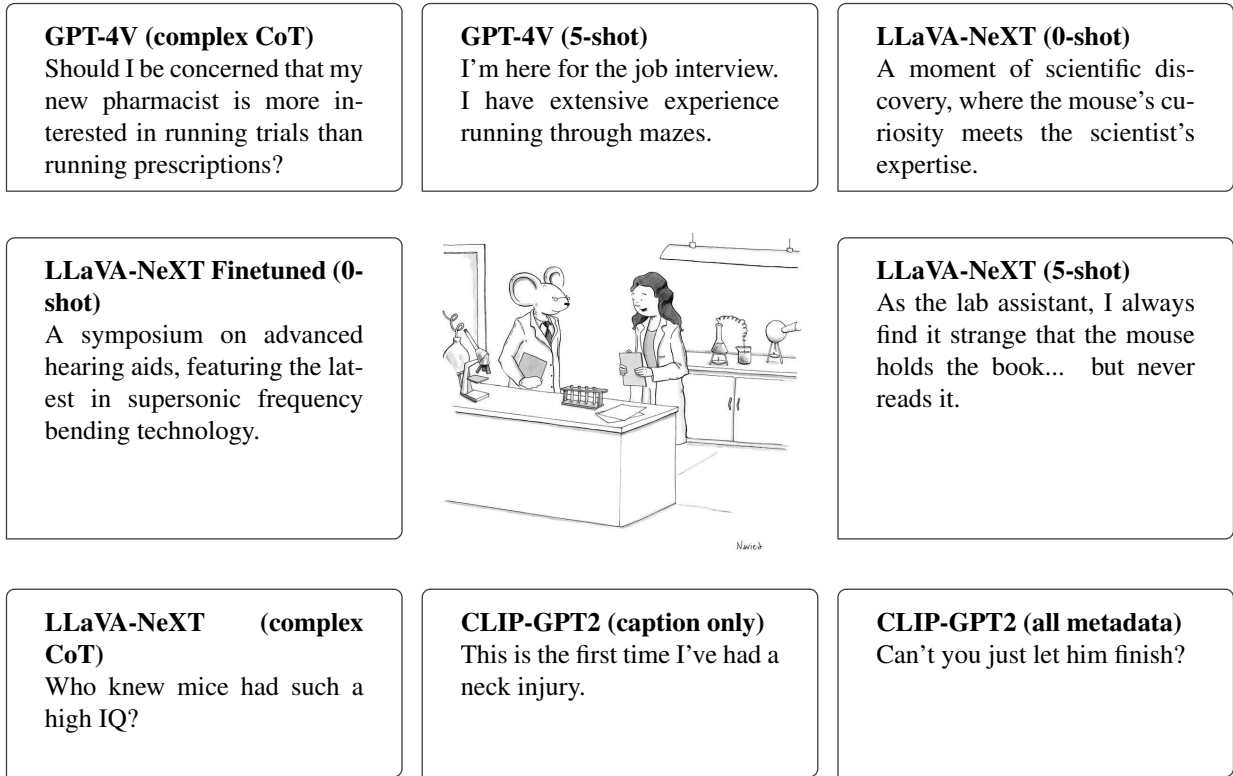


Figure 6: Generated Captions for New Yorker Caption Contest 54.

Winning caption: *And, when you get hungry, the cafeteria is to your right, left, left, right, left, straight, right, straight, left, and then you push on the big lever.*

matic summarizing and machine translation. Figure 7 shows the ROUGE scores for each of our models on the test set. While our models performed poorly with respect to these automated metrics, we will later discuss why these automated metrics, which are designed for machine translation / summarization tasks, do not serve as a good metric for this usecase.

5.2. Manual Quality Evaluation

To surface a more accurate metric for this task, we propose a criteria called *SS-SCORE* metric. For each of the 130 test images, we compare all the machine-generated captions with each other, using the human-written caption as a reference, and select the best machine-generated caption. The quality of the generated caption were based on the following criteria:

1. Captions should contain descriptive information about the cartoon that extends beyond listing features in the image. An observer who examines the image then subsequently internalizes the caption should most often understand the meaning behind why a situation was described the way it was or why a character would have contributed their spe-

cific dialogue.

2. Good captions do not necessarily have to exhaustively describe the situation, but rather convey humor, wit, or a voice that aptly fits within the context of the scene.
3. A caption that addresses similar themes as the winning caption in the associated contest may lead to a more favorable selection criteria on the grounds of resembling an officially winning entry.

After tallying each test example, for each model, we aggregate the statistics by calculating the number of times that model generated a caption that was deemed the best, normalizing by the total number of test examples to obtain a proportion which we equate to the *SS-SCORE*. These metrics can be found in Figure 7.

6. Discussion

6.1. Automated Metrics

From Figure 7, we note that according to the various types of ROUGE scores, the CLIP-GPT2 model performs best, and by prepending the metadata to the

target labels, we achieve a marginally higher ROUGE score. However, automated evaluation metrics such as BLEU and ROUGE do not accurately capture the quality of machine-generated captions. Since BLEU and ROUGE both operate on n -grams, these automated metrics are capturing the commonalities between machine generated and human-written captions at the individual word level without regards to semantic meaning. For a task as sophisticated as this one, word-level similarity is highly uncorrelated with caption quality because there are multiple ways to write humorous captions. Rather than regurgitating the content of the cartoon, contest participants are required to execute complex reasoning that relies on social and cultural nuances related to the cartoon, and based on the diverse range of high quality submissions for a single image, there is no single method for generating a good caption.

6.2. Manual Quality Evaluation

Based on *SS-SCORE*, we notice that GPT-4V in the 5-shot setting performs the best, with GPT-4V in the Chain-of-Thought setting performing second best. This suggests that larger parameter models are able to perform better on this task, likely because they are trained on more data and have a larger knowledge base. GPT-4 models are constantly finetuned on the most recent data, which makes them particularly knowledgeable about both historic and current events; this knowledge is particularly valuable for this cartoon captioning contest, since cartoons from The New Yorker are often timely caricatures of socio-cultural situations.

It is particularly interesting that the 5-shot setting performs almost twice as well as the Chain-of-Thought setting. This suggests that GPT-4V is capable of few-shot learning, and presenting examples of cartoon-caption pairs to GPT-4V appears more helpful for the model to adapt to the specific tone typically found in high quality cartoon captions compared to Chain-of-Thought prompting. Additionally, GPT-4V demonstrated other unique ways of expressing creative generated captions such as the use of puns and double entendres (e.g. a “full-blown service” for a tuba-playing waiter) or references to movies or literature (e.g. “Cloudy with a Chance of Armageddon” for a cloud with a man’s face referring to the movie *Cloudy with a Chance of Meatballs*).

6.3. Qualitative Evaluation

Figure 6 shows one particular test example with the generated captions from all of our models. We also showcase the winning caption selected by The New Yorker magazine.

We note that although the CLIP-GPT2 models are able to mimic the correct tone and style of high quality

captions, they are unable to capture the semantic content of the image. This is likely because the CLIP vision model was trained on natural images, and these cartoons do not exhibit the same color distribution as those natural images. We found that even with fine-tuning, CLIP is unable to capture the proper image features for GPT2 to decode. Moreover, the CLIP-GPT2 encoder-decoder model is far from the state-of-the-art vision transformer models, and due to the relatively small size of the model, it is highly unlikely that it will have the vast social knowledge to create humorous captions.

We also note that LLaVA-NeXT in the 0-shot setting has a tendency to create a *title* for the cartoon rather than a caption for it. Even when we finetune LLaVA-NeXT on the dataset, the model still has the tendency to generate a title. We hypothesize that this is because LLaVA-NeXT does not understand the prompt, and it has no prior knowledge of The New Yorker cartoon caption contest. Thus, it is likely that “New Yorker” was tokenized into an unknown token which has zero knowledge about the magazine and its history, or at the very most has only surface-level knowledge of that linguistic entity.

However, in the 5-shot setting LLaVA-NeXT is able to generate a caption that more closely mimics what could be found in a high quality caption. It begins to use pronouns like in the sentence “I always find it strange...” which indicates that it is learning to understand that high quality captions are typically statements made by characters in the cartoon. Moreover, the Chain-of-Thought setting shows further improvements, with the generated caption demonstrating a clear understanding that character dialogue is important for humorous captions.

Finally, the GPT-4V models perform the best. As discussed before, GPT-4V models have the most recent and largest knowledge database, and are able to successfully incorporate that knowledge into the caption. This is exhibited in its generated captions in Figure 6, which not only shows its understanding of the style of the caption, but also an understanding of what mice typically do in laboratory settings. Thus, we believe that GPT-4V is the current state-of-the-art model for this cartoon captioning task.

7. Future Steps

The task of generating humorous and contextually relevant captions for New Yorker cartoons presents unique challenges that require sophisticated AI models. While our current work demonstrates significant advancements, several avenues for future exploration could enhance the capabilities and performance of such models. Our experiments with models like GPT-4V and LLaVA-NeXT highlight the importance of model size

	ROUGE-1	ROUGE-2	ROUGE-L	SS-SCORE
GPT-4V (CoT)	0.0653	0.0049	0.0591	0.19
GPT-4V (5-shot)	0.0648	0.0019	0.0598	0.37
LLaVA-NeXT (0-shot)	0.0609	0.0006	0.0568	0.09
LLaVA-NeXT Finetuned (0-shot)	0.0637	0.0036	0.0575	0.05
LLaVA-NeXT (5-shot)	0.0586	0.0024	0.0540	0.01
LLaVA-NeXT (CoT)	0.0672	0.0022	0.0617	0.04
CLIP-GPT2 (caption only)	0.0733	0.0062	0.0665	0.11
CLIP-GPT2 (all metadata)	0.0735	0.0039	0.0673	0.14

Figure 7: Evaluation metrics for all models

and the breadth of pre-training data. Future work could involve training even larger models, which inherently possess a more extensive knowledge base and greater capacity for understanding nuanced contexts. Utilizing models like the forthcoming GPT-5 could potentially offer improvements in generating high-quality, humorous captions. Additionally, leveraging distributed training on multiple GPUs or TPUs can facilitate the training of these larger models more efficiently.

Prompt engineering plays a crucial role in guiding the model to produce desirable outputs. Our use of few-shot learning and Chain-of-Thought (CoT) prompting has shown promising results, but further refinement and innovation in prompt design could yield even better performance. Especially for the LLaVA-NeXT model, we noticed many responses misinterpreted the task at hand by describing the image like a traditional caption more often than generating a line that has character or wit. Another option to make the captions more humorous and relevant is to explore hierarchical prompting, where prompts are dynamically adjusted based on intermediate model responses.

While the New Yorker Caption Contest dataset provides a robust foundation, curating additional datasets specifically aimed at caption generation could offer substantial benefits. These datasets could include a broader range of cartoons and comics, along with detailed annotations on humor, context, and cultural references. Collaborating with cartoonists and humorists to annotate datasets with insights into why certain captions are humorous could provide invaluable training data for models.

8. Contributions and Acknowledgment

All people in this project split the project evenly. Stanley worked on training the models, and both team members worked on the data analysis. Sonny worked on formatting the paper and presenting manually evaluating the results.

References

- [1] A. M. Ahmed, M. Yousef, K. F. Hussain, and Y. B. Mahdy. Enhancing image captioning with depth information using a transformer-based framework, 2023.
- [2] A. Chandrasekaran, A. K. Vijayakumar, S. Antol, M. Bansal, D. Batra, C. L. Zitnick, and D. Parikh. We are humor beings: Understanding and predicting visual humor, 2016.
- [3] T. Dettmers, A. Pagnoni, A. Holtzman, and L. Zettlemoyer. QLoRA: Efficient finetuning of quantized LLMs. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [4] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2021.
- [5] J. Hessel, A. Marasović, J. D. Hwang, L. Lee, J. Da, R. Zellers, R. Mankoff, and Y. Choi. Do androids laugh at electric sheep? humor "understanding" benchmarks from the new yorker caption contest, 2023.
- [6] E. J. Hu, yelong shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022.
- [7] L. Jain, K. Jamieson, R. Mankoff, R. Nowak, and S. Sievert. The New Yorker cartoon caption contest dataset, 2020.
- [8] Y. Li, Y. Yu, C. Liang, N. Karampatziakis, P. He, W. Chen, and T. Zhao. Loftq: LoRA-fine-tuning-aware quantization for large language models. In *The Twelfth International Conference on Learning Representations*, 2024.
- [9] C.-Y. Lin and F. J. Och. ORANGE: a method for evaluating automatic evaluation metrics for machine translation. In *COLING 2004: Proceedings of the 20th International Conference on Computational Linguistics*, pages 501–507, Geneva, Switzerland, aug 23–aug 27 2004. COLING.
- [10] T.-Y. Lin, M. Maire, S. Belongie, L. Bourdev, R. Girshick, J. Hays, P. Perona, D. Ramanan, C. L. Zitnick, and P. Dollár. Microsoft coco: Common objects in context, 2015.

- [11] H. Liu, C. Li, Y. Li, and Y. J. Lee. Improved baselines with visual instruction tuning, 2023.
- [12] H. Liu, C. Li, Y. Li, B. Li, Y. Zhang, S. Shen, and Y. J. Lee. Llava-next: Improved reasoning, ocr, and world knowledge, January 2024.
- [13] H. Liu, C. Li, Q. Wu, and Y. J. Lee. Visual instruction tuning. *ArXiv*, abs/2304.08485, 2023.
- [14] K. Nakashima, H. Kataoka, A. Matsumoto, K. Iwata, and N. Inoue. Can vision transformers learn without natural images?, 2021.
- [15] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics, ACL '02*, page 311–318, USA, 2002. Association for Computational Linguistics.
- [16] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. Pytorch: An imperative style, high-performance deep learning library. *ArXiv*, abs/1912.01703, 2019.
- [17] W. Qin, R. Xu, P. Huang, X. Wu, H. Zhang, and L. Luo. What a whole slide image can tell? subtype-guided masked transformer for pathological image captioning, 2023.
- [18] D. Radev, A. Stent, J. Tetreault, A. Pappu, A. Iliakopoulou, A. Chanfreau, P. de Juan, J. Vallmitjana, A. Jaimes, R. Jha, and R. Mankoff. Humor in collective discourse: Unsupervised funniness detection in the New Yorker cartoon caption contest. In *LREC*, 2016.
- [19] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever. Learning transferable visual models from natural language supervision, 2021.
- [20] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever. Language models are unsupervised multi-task learners. 2019.
- [21] D. Shahaf, E. Horvitz, and R. Mankoff. Inside jokes: Identifying humorous cartoon captions. In *KDD*, 2015.
- [22] D. Sun, Y. Bao, and X. Cao. A lightweight transformer for remote sensing image change captioning, 2024.
- [23] A. Teramoto, A. Michiba, Y. Kiriya, T. Tsukamoto, K. Imaizumi, and H. Fujita. Automated report generation for lung cytological images using a cnn vision classifier and multiple-transformer text decoders: Preliminary study, 2024.
- [24] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need, 2023.
- [25] R. Vedantam, X. Lin, T. Batra, C. L. Zitnick, and D. Parikh. Learning common sense through visual abstraction. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 2542–2550, 2015.
- [26] J. Wei, X. Wang, D. Schuurmans, M. Bosma, b. ichter, F. Xia, E. Chi, Q. V. Le, and D. Zhou. Chain-of-thought prompting elicits reasoning in large language models. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837. Curran Associates, Inc., 2022.
- [27] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. Le Scao, S. Gugger, M. Drame, Q. Lhoest, and A. Rush. Transformers: State-of-the-art natural language processing. In Q. Liu and D. Schlangen, editors, *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online, Oct. 2020. Association for Computational Linguistics.
- [28] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhutdinov, R. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention, 2016.
- [29] C. L. Zitnick, D. Parikh, and L. Vanderwende. Learning the visual interpretation of sentences. In *2013 IEEE International Conference on Computer Vision*, pages 1681–1688, 2013.