

This WACV workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

NuScenes-MQA: Integrated Evaluation of Captions and QA for Autonomous Driving Datasets using Markup Annotations

Yuichi Inoue, Yuki Yada, Kotaro Tanahashi, Yu Yamaguchi Turing Inc.

{y.inoue, yu.yamaguchi}@turing-motors.com

Abstract

Visual Question Answering (VQA) is one of the most important tasks in autonomous driving, which requires accurate recognition and complex situation evaluations. However, datasets annotated in a QA format, which guarantees precise language generation and scene recognition from driving scenes, have not been established yet. In this work, we introduce Markup-QA, a novel dataset annotation technique in which QAs are enclosed within markups. This approach facilitates the simultaneous evaluation of a model's capabilities in sentence generation and VQA. Moreover, using this annotation methodology, we designed the NuScenes-MQA dataset. This dataset empowers the development of vision language models, especially for autonomous driving tasks, by focusing on both descriptive capabilities and precise QA. The dataset is available at https://github.com/turingmotors/NuScenes-MQA.

1. Introduction

Generating accurate text explanations for visual scenes has become crucial in the progress of Vision Language Models (VLMs) integrated with Large Language Models (LLMs). It is believed that many challenges can be more effectively tackled by describing visuals in natural language using VLMs. Specifically, autonomous driving requires accurate recognition and complex situation evaluations. This context underscores the potential benefits of VLMs, which can harness the superior logical reasoning capabilities of LLM. As a result, the pursuit of VLMs tailored for autonomous driving is currently a topic of intense research interest [15].

The trend of using autonomous driving datasets enriched with natural language annotations has been growing steadily [11, 16, 17, 28, 34, 41]. By using the data described in natural language, we can align with LLMs that possess general knowledge and high logical reasoning capabilities. By incorporating LLMs, there is potential to create a more intelligent autonomous driving system. To achieve this, it might be necessary to formulate a VLM and curate a relevant dataset for its training. However, datasets annotated in a QA format, which ensures precise language generation and scene recognition from driving scenes, are still a challenge to obtain.

Interpreting visual content accurately is not only essential for autonomous driving research, but is also crucial across diverse tasks. Among them, VQA, which provides accurate descriptions of the visual scenes, is particularly important. Various VQA datasets have been proposed, continuously demonstrating their significant value in training and evaluating state-of-the-art VLMs. Traditional QA tasks have predominantly focused on predicting a singular word. However, with the recent proliferation of sophisticated high-performance LLMs [7], predicting just one word may not fully harness their potential and may even suppress their inherent linguistic generative abilities. To holistically validate the model's comprehension of the visual content, LLaVA-RLHF [37] employed RLHF (Reinforcement Learning from Human Feedback) to counteract the vision language model's hallucination. However, this proved to be both time-consuming and costly.

To address this, we introduced "Markup-QA", wherein the QA segment within a naturally composed text is enclosed by our unique markups. Post-processing extracts this markup-wrapped segment, enabling the evaluation of the accuracy of QAs embedded in the text. By removing the markup, the text retains its completeness, allowing us to assess the model's text generation capabilities using standard evaluation metrics. Another advantage of Markup-QA is its flexibility, allowing for the extraction of QAs from any text segment and embedding multiple QAs within a single sentence, marking an innovative departure in QA tasks.

Using the rich annotations of nuScenes [8] concerning spatial object information, we systematically generated natural language annotations embedded with Markup-QA in a rule-based manner. This dataset, named *NuScenes-MQA*, comprises 1, 459, 933 annotations, covering aspects such as object presence, counts, proximity, and relative positions.



Question: How many <obj>cars</obj> are in <cam>front</cam> of the ego car? Answer: There are <target><cnt>2</cnt> <obj>cars</obj></target>.



Question: Can you specify the exact location of the closest objects and its category? Answer: The closest object to the ego-car is a <obj-car</obj> located at coordinates <0c>(4.24, -9.08)</0c>



Question: What is the object in <cam>front right</cam> of the ego car? Answer: There are <target>cnth=1</cnt> <obj>truck</obj></target> and <target>cnth=1</cnt> <obj>traffic cone</obj></target> in the front right



Question: What is the closest object to the ego vehicle? Answer: The closest object to the ego vehicle is the <obj>traffic cone</obj> with a distance of <ds>3.66</ds>

Figure 1. Example Scenes and Annotations from NuScenes-MQA Dataset

Using *NuScenes-MQA* ensures simultaneous evaluations of accurate QA capabilities and natural language proficiency. Our contributions can be summarized as follows:

- We introduced Markup-QA, a novel dataset annotation technique in which QAs are enclosed within markups. By using data annotated with Markup-QA, QA tasks can be embedded within natural sentences, allowing a concurrent evaluation of textual quality and QA accuracy.
- We proposed and publicly released the NuScenes-MQA dataset annotated in the Markup-QA style, along with its evaluation methodology.
- Using VLMs capable of handling multiple images, we established a baseline for the NuScenes-MQA dataset.

2. Related Work

2.1. Vision Language Datasets in Driving Scenes

Datasets for autonomous driving are inherently multimodal, curated from a variety of sensors [8, 9, 36, 42]. Recently, several existing autonomous driving datasets collected from these sensors have been augmented with textual annotations. For example, the BDD-X [20] supplements driving conditions with textual descriptions that explain the underlying reasons. In addition to describing driving scenarios, DriveGPT4 [41] uses off-the-shelf object detection models in conjunction with GPT, improving BDD-X with recognition tasks and textual captions. The Honda DRAMA dataset [27] introduces the challenge of localizing risk objects and explicating their risks. In pursuing recognitionspecific datasets, NuScenes-QA [33] utilizes pre-annotated object information to propose a 3D VQA task. Moreover, DriveLM [14] constructs a dataset for nuScenes that encapsulates perception, prediction, and planning, all described in text. The trend underscores the growing attention towards leveraging rich sensor data in autonomous driving to recognize spatial object information and articulate it through natural language.

2.2. VQA

The task of VQA involves processing an image and a natural language question to produce a concise natural language response. A wide variety of datasets, such as VQA [4], VQA v2.0 [18], GOA [19], and Visual Genome [21], have been introduced. In particular, in the field of autonomous driving, NuScenes-QA [33], which incorporates position information from surrounding objects, is noteworthy. Early research predominantly combined CNNbased image feature extractor with RNNs [3, 26, 29]. However, with the emergence of Transformer architecture [39], transformer-based language models have become the dominant choice for performance. This trend is evident with the introduction of models such as the encoder-decoder-based PALI [1, 12, 13], decoder-only LLaVA [24, 25] and Mini-GPT4 [10, 44], BLIP-2 [22] with the resampling qformer module, and Flamingo [2] with gated cross-attentions.

2.3. Special Words in Vision Language Prompts

Special words improve the efficiency of vision language tasks by better linking visual information to text. The <image> token, for instance, ties image data directly to specific sentences [2,5,32]. QwenVL [5] further innovates using tokens such as <box> and <ref> to address the visual grounding task, specifying which textual phrases correspond to the regions of the image. Kosmos-2 [32] expands on this by associating multiple bounding boxes with phrases using tokens like <box>, <delim>, and location <loc> tokens. These tokens enrich prompts by conveying details not captured by natural language alone.

3. Proposed Dataset

We introduce a novel QA dataset, called NuScenes-MQA, based on nuScenes [8]. Moving away from the conventional short-answer paradigm, our method emphasizes full-sentence responses, enriching both the content and structure of the answers. The dataset targets key aspects of autonomous driving, such as the presence of objects and relative positioning. Using unique markups, we can highlight and evaluate specific information within the answers. With this methodology, we generated 1, 459, 933 QA pairs derived from 34, 149 driving scenarios. Example annotations are shown in Fig. 1.

3.1. Dataset Construction

In order to create a QA dataset, we employed the annotations provided as ground truth in nuScenes. Contrary to traditional QA datasets that typically structure their response sections with one word, we decided to make our answers as full sentences. To enrich the diversity of our QA templates, we used GPT-4 [30] and crafted 50 expressions per template, ensuring semantic consistency. Human reviewers then curated and adjusted a subset of 20 to 30 from these generated expressions.

Our dataset is based on four core concepts.

- **Specific Object Presence:** Questions that ask about the existence and number of specific objects.
- **Objects in Specific Direction:** Questions asking for the number and category of objects in a specific direction.
- **Relative Distance to Ego Vehicle:** Questions about the relative distance to vehicles. For simplicity, we identify the object closest to the vehicle and its corresponding distance.
- **Relative location to Ego vehicle:** Questions about the location of objects. Similarly, we simplified the task to identify the closest object and its coordinates.

Using the extensive information available in the nuScenes annotations, such as the class and location of recognized objects, and the cameras that capture them, we were able to automate the QA creation process.

3.2. Markup Implementation

Traditional QA evaluation often revolves around predicting a single word, a method that tends to compromise sentence generation capabilities. To address this, we incorporated special markups in our dataset. These markups were differentiated for each QA type as follows:

<target>: Encapsulates <cnt> and <obj>.

<obj>: Represents an object, restricted to a single word.

- <cnt>: Represents a count, restricted to a single word.
- <ans>: Represents a binary response, a single word.

<cam>: Represents one of the six cameras.

<dst>: Represents distance.

<loc>: Represents (x, y) coordinates.

By enveloping the target objects with these markups, we can easily extract the relevant words from the answers. For example:

- In the <cam>back</cam>, <target><cnt>3</cnt> <obj>trucks</obj></target>are detected.
- The closest object to the ego-car is a <obj>car</obj> located at coordinates <loc>(3.43, 1.41)</loc>.

This method allows for the simultaneous evaluation of multiple detections. For example, in the phrase "3 trucks", it is necessary to recognize both the class "trucks" and its count "3". In this particular regard, conventional methods are insufficient. Using our markup methodology, we can simultaneously evaluate both elements. Similarly, this framework facilitates the recognition of two or more classes at the same time. In the example provided above, it is possible to accurately respond to two distinct questions regarding the object category and the location of the object. Hence, the usage of markups enables the design of tasks that can simultaneously answer multiple queries.

3.3. Dataset Statistics

In this section, we discuss the statistics of the NuScenes-MQA dataset. An overview of this dataset is provided through word clouds representing the most frequent terms found within the questions and answers. These visualizations are depicted in Fig. 2 (a) and (b), showing the results for the questions and answers, respectively. The word clouds display a particularly diverse range of phrases, especially in the answers.



Figure 2. Statistics of the NuScenes-MQA dataset

3.3.1 Simple QA Task

Specific Object Presence and Objects in Specific Direction tasks facilitate recognition-based QA from images, essentially assessing the presence and count of objects. In Fig. 2 (c), simple QA tasks seeking binary "Yes" or "No" answers show a slight dominance of "Yes" responses, but the bias does not significantly influence model outcomes. To maintain task feasibility and prevent undue complexity, we excluded questions that required counting more than 20 objects. The distribution of the number of objects to be counted is depicted in Fig. 2 (d).

Our dataset harnesses the <target> markup, enabling the embedding of multiple QAs within a single statement. Fig. 2 (e) illustrates the number of targets encapsulated in a single statement, highlighting tasks that predict counting and category. Although tasks identifying a single target are common, tasks requiring multiple target responses are also prevalent. This variety allows for an extended evaluation beyond typical simple QA.

3.3.2 Distance and position distribution

The nuScenes annotations provide valuable information on the location of objects. Using positional relationships between objects and the ego vehicle, we created QA tasks designated as *Relative Distance to Ego vehicle* and *Relative Location to Ego vehicle*. Drawing inspiration from the range of recognition tasks, such as the occupancy prediction [35, 38], we delimited our focus to objects situated within a 40-meter radius. Fig. 2 (f) and (g), which plot the distance and positional relationship of objects, reveal that most are within a 20-meter range in all directions. To the best of our knowledge, our tasks are the first to require textual answers regarding the spatial information of objects.

4. Methodology

4.1. Model Architecture for Markup-QA

For our Markup-QA tasks, we introduce a model that combines a vision transformer (ViT) and a decoder-only language model via a simple linear module. The efficiency of this architecture is supported by previous works such as GIT [40], LLaVA [24,25], and MiniGPTv2 [10]. The architecture of our model is illustrated in Fig. 3.

The visual input undergoes feature extraction using a ViT pre-trained in CLIP, subsequently extracting patch features from the final layer. Considering the nuScenes dataset accommodates images from six distinct camera orientations, we use six ViTs, each dedicated to extracting features from its corresponding camera's image. The ViT shares the parameters. These extracted features are subsequently added with six trainable positional embeddings, in a manner similar to [40]. The ViT patch features are then projected, using a single layer adapter module, to match the size of the text embeddings, thus formulating the visual embeddings. These visual embeddings, similar to text embeddings, are fed into the language model for both training and inference. During training, a causal mask is applied to text embeddings to negate the influence of future information during self-attention computations. However, this constraint is not imposed on visual embeddings, as shown in Fig. 4.



Output text

Figure 3. Vision Language Model Architecture for NuScenes-MQA



Figure 4. Attention Mask for Vision Language Model

4.1.1 Tokenization

Words used for markup adopt a format rarely seen in conventional text. Based on previous studies [5, 32], these markups were typically incorporated as additional tokens. In the results section, we provide a comparison between tokenization using conventional tokenizers and tokenization that treats these markups as additional tokens.

5. Experiments

5.1. Evaluation Metrics

To evaluate the quality of the generated text, we used n-gram-based standard metrics, notably BLEU-1 [31], BLEU-4 [31], METEOR [6], and ROGUE-1 F-score [23]. QA tasks were evaluated in the accuracy metric, while distance measurement tasks were evaluated using the mean absolute error (MAE).

Due to the inherent nature of sentence generation as an evaluation method, inference requires considerable time. To efficiently assess the model's performance, we chose a smaller subset for the model evaluation. From our test set, we carefully extracted 2,000 samples, ensuring a balanced representation of the tasks. Then, our evaluations were executed on this curated subset.

5.2. Training Details

For the ViT, we utilized OpenAI's ViT large patch14, which is pre-trained using CLIP. We used an image resolution of 224x224 during training and evaluation. The language model was OPT [43], with parameter sizes of 125M, 1.3B, and 6.7B. Both the ViT and the language model were initialized using pre-trained parameters, while the adapter module started with random initialization. All parameters spanning the ViT, the language model, and the adapter module were trained.

Our training data combined random QAs from a given scene, ensuring that all types of question were present in a sample. The maximum sequence length for the text segment, excluded from visual embeddings, was set at 256. AdamW optimizer was used, with a learning schedule outlined by a one-cycle scheduler starting at 1e - 6, peaking at 1e - 4, and dropping back to 1e - 6. The other parameters retained their default values. Training was carried out for 10 epochs using the standard cross-entropy loss. The comprehensive training regimen was orchestrated on 8 Nvidia A100 GPUs or 8 Nvidia H100 GPUs.

5.3. Results

Table 1 shows the performance differences in the Sentence Generation Performance (SGP) and VQA models, according to the size of the model and to the use of markup tokens as additional tokens.

Model		SGP					VQA							
LM	Special token	BLEU-1	BLEU-4	METEOR	ROUGE	Avg. SGP	Yes/No	Cat.	Cat. & Count	RD	Cat. RD	Loc. x	Loc. y	Cat. Loc.
OPT-125M		0.659	0.371	0.679	0.616	0.581	0.795	0.710	0.319	3.546	0.424	5.903	6.420	0.458
OPT-1.3B		0.698	0.404	0.677	0.626	0.601	0.795	0.703	0.315	4.287	0.430	6.306	6.088	0.462
OPT-6.7B		0.674	0.384	0.676	0.624	0.589	0.815	0.700	0.317	4.208	0.426	6.511	6.469	0.460
OPT-125M	1	0.644	0.360	0.455	0.474	0.483	0.820	0.763	0.331	2.644	0.504	6.158	6.408	0.520
OPT-1.3B	1	0.681	0.388	0.439	0.477	0.496	0.800	0.714	0.325	4.221	0.434	6.252	6.090	0.458

Table 1. Results of different model settings on the NuScenes-MQA test set. We evaluated the natural language generation capability (SGP) and VQA. Cat. means the accuracy of the object category. RD and Loc. represent the relative distance to the ego vehicle and the relative location to the ego vehicle, respectively.

5.3.1 Sentence Generation Performance

When evaluating models without the special token, the OPT-1.3B model consistently outperformed its counterparts in most metrics. It achieved the highest scores of BLEU-1, BLEU-4, and ROUGE, with values of 0.698, 0.404, and 0.626, respectively. Although the METEOR scores were comparably high across all three models, the OPT-125M slightly surpassed the others with a score of 0.679. Furthermore, the OPT-1.3B model achieved a peak *Avg. SGP* score of 0.601.

In contrast, using markups as a special token led to a marked decline in METEOR scores across all models. Specifically, the OPT-125M model showed the most significant drop, falling to 0.455. Furthermore, not only ME-TEOR, but other scores also experienced a general decline. Interestingly, for BLEU-1, the OPT-1.3B model with the special token retained a slight advantage, recording a score of 0.681. The *Avg. SGP* also decreased with the incorporation of the special token. These observations suggest that integrating markup as a special token may adversely impact the language generation capabilities of the models.

5.3.2 VQA Performance

Focusing on models without additional tokens, the performance in the *Yes/No* metric remained consistent across all models. In the accuracy of categories, *Cat.*, the OPT-125M model had a slight advantage, achieving a score of 0.710. Interestingly, for *Loc.* x and *Loc.* y, none of the models showed impressive results, indicating that the task was particularly challenging.

When the special token was integrated, the OPT-125M model showed notable improvements in several metrics such as *Yes/No*, *Cat.*, and *Cat.* & *Count*, with scores of 0.820, 0.763, and 0.331, respectively. Both the OPT-125M and OPT-1.3B models showed improvement in the *RD* and *Cat. RD* metrics. Remarkably, the OPT-125M model ex-

hibited significant gains, with the scores for *RD* and *Cat. RD* to 2.644 and 0.504, respectively. The OPT-1.3B model exhibited a slight improvement in most of the VQA metrics upon the addition of markup tokens as a special token.

5.3.3 Quantifying the difficulty of multiple QAs

In our dataset, following the criteria defined under *Objects in Specific Direction*, a single sentence encompasses multiple QA tasks that require identification of both the categories of objects and their counts. Table 2 delineates the accuracy rates based on the number of QAs present in a single sentence (*n*-QA). As the number of QAs increases, tasks become more complex, leading to lower accuracy rates. This decrease is more noticeable in *Cat.* & *count* than in *Cat.*. Given these challenges, there is a pressing need for further research to effectively address multiple QAs in a single sentence.

n-QA	Cat.	Cat. & Count
1	0.88	0.42
2	0.59	0.16
3	0.49	0.14
4	0.53	0.16
5	0.51	0.13
6	0.51	0.10

Table 2. Accuracy based on the number of QAs per sentence

6. Conclusion

In this work, we introduced the NuScenes-MQA dataset, which employs a Markup-QA approach where QA is encapsulated within the text using markup. Through the implementation of the Markup-QA scheme, we established a framework that facilitates the simultaneous evaluation of a model's capabilities in sentence generation and VQA. This dataset empowers the development of vision language models, especially for autonomous driving tasks, by focusing on both descriptive capabilities and precise QA. We have also established a baseline model that provides a starting point to demonstrate the practical value of our approach.

7. Limitation

While our research offers significant insight, it comes with certain constraints worth noting. First, the dataset has been constructed using a rule-based approach. As a consequence, it may lack the rich diversity often inherent in natural language. This limitation could potentially make it less ideal for training larger models, such as OPT-6.7B, due to potential insufficiency. Furthermore, the limited variety of tasks that address spatial information raises concerns. Specifically, it remains uncertain whether the model truly captures expressions pertaining to positional data. These limitations underscore areas for deeper future research. We believe that future work will overcome these challenges and further advance the field.

References

- Sina Ahmadi, Milind Agarwal, and Antonios Anastasopoulos. Pali: A language identification benchmark for persoarabic scripts, 2023. 2
- [2] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob Menick, Sebastian Borgeaud, Andrew Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karen Simonyan. Flamingo: a visual language model for few-shot learning, 2022. 2, 3
- [3] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (*CVPR*), June 2018. 2
- [4] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings* of the IEEE International Conference on Computer Vision (ICCV), December 2015. 2
- [5] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. arXiv preprint arXiv:2308.12966, 2023. 3, 5
- [6] Satanjeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on in-*

trinsic and extrinsic evaluation measures for machine translation and/or summarization, pages 65–72, 2005. 5

- [7] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. Advances in neural information processing systems, 33:1877–1901, 2020. 1
- [8] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 1, 2, 3
- [9] Holger Caesar, Juraj Kabzan, Kok Seang Tan, Whye Kit Fong, Eric M. Wolff, Alex H. Lang, Luke Fletcher, Oscar Beijbom, and Sammy Omari. nuplan: A closed-loop mlbased planning benchmark for autonomous vehicles. *CoRR*, abs/2106.11810, 2021. 2
- [10] Jun Chen, Deyao Zhu, Xiaoqian Shen, Xiang Li, Zechu Liu, Pengchuan Zhang, Raghuraman Krishnamoorthi, Vikas Chandra, Yunyang Xiong, and Mohamed Elhoseiny. Minigpt-v2: large language model as a unified interface for vision-language multi-task learning. arXiv preprint arXiv:2310.09478, 2023. 2, 4
- [11] Long Chen, Oleg Sinavski, Jan Hünermann, Alice Karnsund, Andrew James Willmott, Danny Birch, Daniel Maund, and Jamie Shotton. Driving with llms: Fusing object-level vector modality for explainable autonomous driving. *arXiv preprint arXiv:2310.01957*, 2023. 1
- [12] Xi Chen, Josip Djolonga, Piotr Padlewski, Basil Mustafa, Soravit Changpinyo, Jialin Wu, Carlos Riquelme Ruiz, Sebastian Goodman, Xiao Wang, Yi Tay, Siamak Shakeri, Mostafa Dehghani, Daniel Salz, Mario Lucic, Michael Tschannen, Arsha Nagrani, Hexiang Hu, Mandar Joshi, Bo Pang, Ceslee Montgomery, Paulina Pietrzyk, Marvin Ritter, AJ Piergiovanni, Matthias Minderer, Filip Pavetic, Austin Waters, Gang Li, Ibrahim Alabdulmohsin, Lucas Beyer, Julien Amelot, Kenton Lee, Andreas Peter Steiner, Yang Li, Daniel Keysers, Anurag Arnab, Yuanzhong Xu, Keran Rong, Alexander Kolesnikov, Mojtaba Seyedhosseini, Anelia Angelova, Xiaohua Zhai, Neil Houlsby, and Radu Soricut. Palix: On scaling up a multilingual vision and language model, 2023. 2
- [13] Xi Chen, Xiao Wang, Lucas Beyer, Alexander Kolesnikov, Jialin Wu, Paul Voigtlaender, Basil Mustafa, Sebastian Goodman, Ibrahim Alabdulmohsin, Piotr Padlewski, Daniel Salz, Xi Xiong, Daniel Vlasic, Filip Pavetic, Keran Rong, Tianli Yu, Daniel Keysers, Xiaohua Zhai, and Radu Soricut. Pali-3 vision language models: Smaller, faster, stronger, 2023. 2
- [14] DriveLM Contributors. Drivelm: Drive on language. https://github.com/OpenDriveLab/DriveLM, 2023. 2
- [15] Can Cui, Yunsheng Ma, Xu Cao, Wenqian Ye, Yang Zhou, Kaizhao Liang, Jintai Chen, Juanwu Lu, Zichong Yang, Kuei-Da Liao, Tianren Gao, Erlong Li, Kun Tang, Zhipeng Cao, Tong Zhou, Ao Liu, Xinrui Yan, Shuqi Mei, Jianguo Cao, Ziran Wang, and Chao Zheng. A survey on multimodal

large language models for autonomous driving. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) Workshops*, 2024. 1

- [16] Thierry Deruyttere, Simon Vandenhende, Dusan Grujicic, Luc Van Gool, and Marie-Francine Moens. Talk2car: Taking control of your self-driving car. arXiv preprint arXiv:1909.10838, 2019. 1
- [17] Daocheng Fu, Xin Li, Licheng Wen, Min Dou, Pinlong Cai, Botian Shi, and Yu Qiao. Drive like a human: Rethinking autonomous driving with large language models. arXiv preprint arXiv:2307.07162, 2023. 1
- [18] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 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 (CVPR)*, July 2017. 2
- [19] Drew A. Hudson and Christopher D. Manning. 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* (*CVPR*), June 2019. 2
- [20] Jinkyu Kim, Anna Rohrbach, Trevor Darrell, John Canny, and Zeynep Akata. Textual explanations for self-driving vehicles. *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 2
- [21] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vi*sion, 123(1):32–73, 2017. 2
- [22] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. BLIP-2: bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*, 2023. 2
- [23] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004. 5
- [24] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning, 2023. 2, 4
- [25] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning, 2023. 2, 4
- [26] Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Hierarchical question-image co-attention for visual question answering. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. 2
- [27] Srikanth Malla, Chiho Choi, Isht Dwivedi, Joon Hee Choi, and Jiachen Li. Drama: Joint risk localization and captioning in driving. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1043–1052, January 2023. 2
- [28] Jiageng Mao, Yuxi Qian, Hang Zhao, and Yue Wang. Gpt-driver: Learning to drive with gpt. arXiv preprint arXiv:2310.01415, 2023. 1

- [29] Duy-Kien Nguyen and Takayuki Okatani. Improved fusion of visual and language representations by dense symmetric co-attention for visual question answering. *CoRR*, abs/1804.00775, 2018. 2
- [30] OpenAI. Gpt-4 technical report, 2023. 3
- [31] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318, 2002. 5
- [32] Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu Wei. Kosmos-2: Grounding multimodal large language models to the world, 2023. 3, 5
- [33] Tianwen Qian, Jingjing Chen, Linhai Zhuo, Yang Jiao, and Yu-Gang Jiang. Nuscenes-qa: A multi-modal visual question answering benchmark for autonomous driving scenario. arXiv preprint arXiv:2305.14836, 2023. 2
- [34] Hao Sha, Yao Mu, Yuxuan Jiang, Li Chen, Chenfeng Xu, Ping Luo, Shengbo Eben Li, Masayoshi Tomizuka, Wei Zhan, and Mingyu Ding. Languagempc: Large language models as decision makers for autonomous driving. arXiv preprint arXiv:2310.03026, 2023. 1
- [35] Chonghao Sima, Wenwen Tong, Tai Wang, Li Chen, Silei Wu, Hanming Deng, Yi Gu, Lewei Lu, Ping Luo, Dahua Lin, and Hongyang Li. Scene as occupancy. 2023. 4
- [36] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, Vijay Vasudevan, Wei Han, Jiquan Ngiam, Hang Zhao, Aleksei Timofeev, Scott Ettinger, Maxim Krivokon, Amy Gao, Aditya Joshi, Yu Zhang, Jonathon Shlens, Zhifeng Chen, and Dragomir Anguelov. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 2
- [37] Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liang-Yan Gui, Yu-Xiong Wang, Yiming Yang, Kurt Keutzer, and Trevor Darrell. Aligning large multimodal models with factually augmented RLHF. Sept. 2023. 1
- [38] Xiaoyu Tian, Tao Jiang, Longfei Yun, Yue Wang, Yilun Wang, and Hang Zhao. Occ3d: A large-scale 3d occupancy prediction benchmark for autonomous driving. arXiv preprint arXiv:2304.14365, 2023. 4
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. 2
- [40] Jianfeng Wang, Zhengyuan Yang, Xiaowei Hu, Linjie Li, Kevin Lin, Zhe Gan, Zicheng Liu, Ce Liu, and Lijuan Wang. Git: A generative image-to-text transformer for vision and language. arXiv preprint arXiv:2205.14100, 2022. 4
- [41] Zhenhua Xu, Yujia Zhang, Enze Xie, Zhen Zhao, Yong Guo, Kenneth KY Wong, Zhenguo Li, and Hengshuang Zhao.

Drivegpt4: Interpretable end-to-end autonomous driving via large language model. *arXiv preprint arXiv:2310.01412*, 2023. 1, 2

- [42] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 2
- [43] 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. Opt: Open pre-trained transformer language models, 2022. 5
- [44] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigpt-4: Enhancing vision-language understanding with advanced large language models. arXiv preprint arXiv:2304.10592, 2023. 2