

Human-Centric Autonomous Systems With LLMs for User Command Reasoning

Yi Yang^{1,2}, Qingwen Zhang¹, Ci Li¹, Daniel Simões Marta¹, Nazre Batool², John Folkesson¹
¹ KTH Royal Institute of Technology ² Scania AB

{yiya, qingwen, cil, dlmarta, johnf}@kth.se {carol-yi.yang, nazre.batool}@scania.com

Abstract

The evolution of autonomous driving has made remarkable advancements in recent years, evolving into a tangible reality. However, a human-centric large-scale adoption hinges on meeting a variety of multifaceted requirements. To ensure that the autonomous system meets the user’s intent, it is essential to accurately discern and interpret user commands, especially in complex or emergency situations. To this end, we propose to leverage the reasoning capabilities of Large Language Models (LLMs) to infer system requirements from in-cabin users’ commands. Through a series of experiments that include different LLM models and prompt designs, we explore the few-shot multivariate binary classification accuracy of system requirements from natural language textual commands. We confirm the general ability of LLMs to understand and reason about prompts but underline that their effectiveness is conditioned on the quality of both the LLM model and the design of appropriate sequential prompts. Code and models are public with the link https://github.com/KTH-RPL/DriveCmd_LLM.

1. Introduction

Autonomous driving (AD) has experienced significant advances in recent years [1–3], enhanced by advancements in both hardware and machine learning techniques. The field has seen diverse developments: multi-task learning approaches [4] over traditional modular systems [5], end-to-end methodologies [6, 7] with planning-oriented techniques [8, 9], and widespread adoption of Bird-Eye-View representations [10]. Many explorations have been conducted towards safe and robust autonomous systems.

Nonetheless, several aspects are often overlooked, such as taking into consideration human intent when designing AD systems [11]. The future of AD should integrate human-centered design with advanced AI capabilities to reason and interpret the user’s intent [12, 13]. This integration offers numerous advantages e.g. language interaction [14], driving scene understanding [15], contextual rea-

soning [16], as well as explainability and trust. In a recent milestone survey, Chen et al. [11] emphasize the necessity of considering human behaviors and AD systems to ensure communication transparency and efficiency. They also hint towards a human-machine hybrid intelligence approach, asserting that the reliability of intelligent systems [17] depends on incorporating learnability and the influence of human intelligence and mentorship.

A current relevant challenge stems from allowing humans to interact using natural language with AD systems [15, 18], aiming to adhere to both human preferences and intent. Our contribution here is a series of insights and experiments on a variety of LLM models and prompt designs exploring the interaction between users and AD systems via verbal commands.

While the use of natural language inherently poses a challenge in itself, we leverage recent advancements in large pre-trained foundational models [19–21] alongside their zero-shot capabilities to adapt to new tasks [22]. In our approach, we adopt a few-shot learning strategy, keeping an LLM frozen and enhancing reasoning through sequential prompting, drawing inspiration from previous works [12, 13, 23]. Ultimately, our goal is to condition an LLM for the task of multivariate binary classification of AD systems based on in-cabin user commands.

2. Related Work

Few-shot Reasoning Capabilities of LLMs. Recent research indicates that LLMs with sufficient expansiveness possess the ability to perform sophisticated reasoning tasks [22, 24]. In this work, we adopt an approach somewhat akin to Socratic models [22], where models—encompassing vision, language, and sound—can operate in a zero-shot or few-shot fashion. A prevalent strategy in this type of approach involves keeping certain model subcomponents static, particularly those relevant to a single modality, preserving them for use in subsequent tasks [25, 26]. This methodology interchangeably aligns well with few-shot transfer learning [27–30], wherein an LLM, originally trained on a vast array of internet-scale text prompts for various tasks such as text completion and sentiment

LLM prompt conditioning and prompt design with user examples

| | |
|---------------------|--|
| a) LLM Conditioning | <p>You'll receive a command message for a self-driving vehicle. Follow these steps to respond:</p> <p>Step 1: First decide whether the external perception system is required for this command. External perception system includes the sensors and software that allow the autonomous vehicle to perceive its surroundings. It typically includes cameras, lidar, radar, and other sensors to detect objects, pedestrians, other vehicles, road conditions, and traffic signs/signals. For example, any movement, sense or detect the surroundings.</p> <p>Step 2: Answer "Is in-cabin monitoring required?" in-cabin monitoring involves cameras, thermometers, or other sensors placed inside the vehicle's cabin to monitor the state of occupants and other conditions. It includes everything in-cable system, for example, seats, windows, doors, multimedia system, alert system, etc.</p> <p>... ..</p> <p>Step 8: Answer "Is there a possibility of violating traffic laws?" Violating traffic laws refers to any action performed by the vehicle that goes against the established traffic regulations of the region. An autonomous vehicle's system is typically designed to adhere strictly to traffic laws. It includes anything risk command. For example, related to the traffic laws, like speed, traffic light, emergency action etc.</p> <p>Answer the 8 questions use the following format:</p> <p>Step 1: 'Yes' or 'No' <step 1 reasoning></p> <p>Step 2: 'Yes' or 'No' <step 2 reasoning></p> <p>... ..</p> <p>Step 8: 'Yes' or 'No' <step 8 reasoning></p> <p>Response to user: Output is [A1 A2 A3 A4 A5 A6 A7 A8]. Replace A1-A8 with 1 for 'Yes' and 0 for 'No'.</p> |
| b) User Assistant | <p>[Provide few-shots for LLM to learn. Show one shot as an example below.]</p> <p>Call my friend Carol.</p> <p>Explanation:</p> <p>Step 1: No, it doesn't involve any movement, sense or detect the surrounding.</p> <p>Step 2: Yes, it requires to use the in-cabin multimedia to call the people.</p> <p>... ..</p> <p>Step 8: No, it is not a risky command.</p> <p>Therefore, the output should be : [0 1 0 0 1 1 1 0].</p> |
| c) User GPT-4 [21] | <p>[Test on real commands.]</p> <p>Make an illegal U-turn here.</p> <p>Explanation:</p> <p>Step 1: Yes, it needs to sense or detect the surrounding for making a u-turn.</p> <p>Step 2: No, it doesn't involve anything inside the vehicle's in-cabin system, like seats, windows, doors, multimedia system, alert system, etc.</p> <p>... ..</p> <p>Step 8: Yes, it is a risky command and it is illegal to make a u-turn</p> <p>Therefore, the output should be : [1 0 1 1 0 0 0 1].</p> |

Table 1. Macroscopic view of our approach. For visual readability, steps are in orange (each step contains one question), user inputs are in purple, explanations are in blue, and the corresponding LLM classification output is in green. a) LLM initial conditioning of several steps in the form of a prompt. b) Few-shot examples resembling a *chain-of-thought* [23] to improve reasoning. c) Example with test commands on GPT-4.

analysis, is repurposed for a new, specific task. In our case, this involves leveraging general domain knowledge of AD for a classification challenge [31]. However, the success of adapting language models to new tasks largely depends on the prompting strategy and the quality of the prompts [32]. Wei et al. [23] demonstrated that standard prompting is often insufficient. They introduced a *chain-of-thought* prompting strategy, where a sequence of thoughtful demonstrations significantly improves performance in tasks such as symbolic reasoning, commonsense, and arithmetic. Inspired by the above works, we aim to keep the weights of an LLM frozen and design a chain-of-thought strategy for a downstream task of classification from in-cabin user commands.

Human-centric Autonomous Driving. Incorporating human intent in the form of natural language within AD is a relatively new and active field of research. The emergence of recent datasets and works has significantly boosted its development. Datasets like NuPrompt [33], NuScenes-QA [34], DRAMA [35], enhance the autonomous driving datasets with provided texts, for different tasks including object tracking, visual question answering, image caption and so on. As many works have shown promising results, a possible goal would be to mediate planning-oriented AD [8], which includes perception, prediction and planning, with human intent. In recent exploratory work, there is verified evidence for the effectiveness of integrating LLMs into human-centric AD [36–39]. In a survey by Zhou et

| Methods | Accuracy (%) ↑ | | | | | | | | | |
|------------------------|----------------|----------------|------------|---------------------|--------------|-----------------|---------------|---------------|----------------|--------------|
| | Command Level | Question Level | | | | | | | | |
| | | Overall | Perception | In-cabin Monitoring | Localization | Vehicle control | Entertainment | Personal Data | Network Access | Traffic laws |
| Random | 0.36 | 49.44 | 50.59 | 47.68 | 51.96 | 48.95 | 47.86 | 48.23 | 50.59 | 49.68 |
| Rule-based* | 4.09 | 69.19 | 62.33 | 74.52 | 69.70 | 65.15 | 87.99 | 70.43 | 58.60 | 64.79 |
| CodeLlama-34b-Instruct | 16.28 | 74.56 | 65.61 | 68.61 | 78.53 | 79.71 | 75.98 | 70.15 | 75.07 | 82.80 |
| Llama-2-70b-Chat | 21.29 | 82.46 | 87.17 | 83.44 | 90.35 | 87.81 | 79.98 | 77.07 | 85.90 | 67.97 |
| GPT-3.5 | <u>36.21</u> | <u>88.03</u> | 88.63 | 86.44 | 90.17 | 87.44 | 96.09 | 80.98 | 87.90 | 86.62 |
| GPT-4 | 38.03 | 89.02 | 93.18 | 74.89 | 91.54 | 88.63 | 94.45 | 85.99 | 91.99 | 91.54 |

Table 2. Performance comparison of different methods or models on the user command benchmark regarding 8 yes/no questions. We compare with the methods of random guess and rule-based as baseline. *Note that the rules are generated by the GPT-4 with Advanced Data Analysis plug-in. With our designed prompt shown in Tab. 1, we compare several LLMs including Llama, CodeLlama, GPT-3.5-turbo, and GPT-4. The best results are shown in **bold** and the second best results are shown in underlined.

al. [40], large language and vision models show potential to contribute to the AD system in different submodules including perception and understanding, navigation and planning, decision-making and control, as well as end-to-end pipelines. For instance, as explored by Fu et al. [18], understanding the common sense driving intentions embedded in human commands is achieved by utilizing the reasoning capacities of LLMs in addressing long-tailed cases. This type of approach could be instrumental in the evolution of AD technologies that mirror human-like driving nuances. Ding et al. [41] use multimodal LLM to inform AD system the localization of risk objects and provide suggestions for safety and robustness. More closely related to our work is that reported by Cui et al. [12, 13] and Jain et al. [42], demonstrating that utilizing linguistics and understanding from verbal commands can improve driving decisions and enhance personalized driving experiences through ongoing verbal feedback. In this exploratory work, we aim to holistically address some of the weaknesses posed when handling long-tail data or out-of-distribution driving scenarios by tapping into the overall implicitly acquired domain knowledge of LLMs [31, 43].

3. Method

Given user commands as input text sentences, the model is expected, through reasoning and sentiment analysis, to provide binary results (‘Yes’ / ‘No’) for 8 specific questions. These questions inquire about the command’s diverse requirements- ranging from perception, in-cabin monitoring, localization, control, network access, and entertainment to human privacy and traffic laws.

Since LLMs undergo training on extensive datasets, integrating them within autonomous systems significantly enhances AD systems’ capability to grasp both scene dynamics and user intent. In this paper, we utilize both online (from the GPT series [21]) and offline (from the Llama se-

ries [44]) LLMs. As demonstrated in Tab. 1, we begin by informing the LLM that its task is to engage in a question and answer (Q&A) format by responding with ‘Yes’ or ‘No’. By instructing the model with an in-depth explanation for each of the eight questions, it is prompted to process information step by step. To further optimize accuracy, the LLMs are provided with few examples of in-context few-shot learning [20]. Finally, conditioned on the real user command, LLMs are instructed to produce both step-by-step explanations and results in a predetermined format.

4. Experiment

Datasets and Metrics: We assess the performance of LLMs on 1,099 in-cabin user commands from the UCU Dataset [45]. The details are available on the challenge website¹. The evaluation involves determining if a command requires any of the eight specified modules’ help to achieve autonomy. Official evaluation metrics include accuracy at *the question level* (accuracy for each individual question) and at *the command level* (accuracy is only acknowledged if all questions for a particular command are correctly identified).

LLM Models and Baseline: We benchmark our prompt approaches using a range of LLMs, including GPT-3.5-turbo / GPT-4², and CodeLlama-34b-Instruct / Llama-2-70b-Chat [44]. The GPT models are accessed online, whereas Llama models offer an on-board solution, with both their code and pretrained weights open-sourced. CodeLlama-34b-Instruct is an enhanced version of the original Llama [46] additionally trained on code generation and instruction problems with 34 billion parameters in the model. Llama-2-70b-Chat has roughly double the number

¹<https://llvm-ad.github.io/challenges/>

²<https://platform.openai.com>. Note that during the time we used the models, GPT-3.5-turbo pointed to GPT-3.5-turbo-0613 version, and GPT-4 pointed to GPT-4-0613 version.

| Detailed Explanation | | Few Shots | Level-based Accuracy (%) \uparrow | |
|----------------------|--------|-----------|-------------------------------------|--------------|
| w/o step | w step | | Command | Question |
| ✓ | ✓ | ✓ | 13.00 | 75.00 |
| | | | 22.00 | 81.00 |
| | | | 26.00 | 79.00 |
| ✓ | ✓ | ✓ | 35.00 | 87.00 |
| | | | 42.00 | 88.63 |
| | | | <u>40.00</u> | <u>87.88</u> |

Table 3. Ablation study in the impact of different prompts.

| # of shot | Level-based Accuracy (%) \uparrow | |
|-----------|-------------------------------------|--------------|
| | Command | Question |
| 0 | 26.00 | 79.00 |
| 1 | 23.00 | 80.63 |
| 2 | <u>40.00</u> | <u>87.63</u> |
| 3 | 41.00 | 87.88 |
| 4 | <u>40.00</u> | 87.88 |

Table 4. Ablation study in the number of provided shots.

of parameters and is trained mainly on conversational interactions enhanced with human reinforcement learning.

We compare the LLMs’ performance against two baselines. The first employs a simple random guessing strategy. The second is a rule-based method that identifies specific keywords in the user command to determine system requirements. It is noteworthy that these rules are automatically generated by ChatGPT-4, augmented with the Advanced Data Analysis plug-in.

Evaluation: The results are presented in Tab. 2. Compared to the random guess and rule-based approaches, LLMs exhibit notably higher accuracy, especially at the command level. Among the LLMs tested, the GPT series surpasses the Llama models, where GPT-4 achieves a peak accuracy of 89.02% at the question level and 38.03% at the command level. Diving into question-specific accuracy, GPT consistently behaves excellently in determining if a command requires perception, localization, entertainment submodules, or if it might violate traffic laws. However, for the in-cabin monitoring, GPT’s predictions display some inconsistency and lower accuracy, especially for GPT-4 (74.89%). It is observed that in the evaluation of 1,099 commands, GPT-4 incorrectly responded to 276 queries related to in-cabin monitoring. A majority proportion of these errors, amounting to 259 instances, involved GPT-4 saying “Yes, it requires in-cabin monitoring” when the ground truth doesn’t, with only a few errors being the reverse. To further investigate why GPT-4 says yes often, 2 keywords are found in its explanations: *multimedia* and *alert* system. Both terms are typically associated with in-cabin monitor-

ing in our given instructed prompts. Specifically, 155 commands (56%) are related to multimedia - GPT-4 identifies these as requiring in-cabin multimedia activities, such as making calls, playing radio or videos, or screen displays. Additionally, in 26% of the instances (71 commands), GPT-4 associates the command with the utilization of in-cabin monitoring systems for alerting or notifying the vehicle’s occupants. Given this context, it’s difficult to conclusively label GPT-4’s responses as incorrect. This is largely due to the vague definition of what exactly constitutes in-cabin monitoring in the context of our study.

Ablation Study: We perform two ablation studies to explore the impacts of providing a detailed explanation of instructions and the given number of shot examples. To reduce costs, we run experiments on GPT-3.5 with a small test subset that matches the distribution of the original test data created by ChatGPT-4. More specifically, ChatGPT-4 is asked to analyze the data distribution for each question and sample 100 commands according to it.

The performance of combining three prompting methods is shown in Tab. 3. Note that for providing a detailed instruction explanation, we experimented with two formats: a step-by-step approach as illustrated in Tab. 1, and a consolidated paragraph format (labeled as ‘w/o step’ in Tab. 3). For the latter case, we omit the word `step: #` which causes less structured prompts. Each prompting method exhibits enhanced performance, with the few-shot method outperforming two detailed explanation methods. Notably, combining one of the prompting methods from detailed explanation and the few-shot method significantly augments the performance.

Tab. 4 shows the impact of the number of given few shot examples with the step-by-step detailed explanation method. We observe a notable increase when two examples are provided, which indicates the importance of providing shot examples for performance gain. However, as the number of examples continues to increase, the impact on the final results becomes subtle, suggesting the existence of a saturation threshold or limitations.

5. Conclusion

In this paper, we offered a human-centric perspective by providing several key insights with different prompting designs to condition LLMs towards achieving AD system requirements from verbal user commands. Our work included an analysis of several online and offline popular LLMs with a variety of ablation studies. We hope that this work inspires our community to consider human intent in the design of an AD system. Future work points towards the direction of incorporating human feedback in a more intelligent and effective way. This would pave the way for the development of AD systems that are not only more reliable and capable in their reasoning and comprehension skills

but also more closely aligned with human-like preferences and behaviors. The goal is to establish an AD system with human-centric values, thereby fostering greater trust and acceptance among users.

Acknowledgement

Thanks to RPL’s members: Xuejiao Zhao and Boyue Jiang, who gave constructive comments on this work. We also thank the anonymous reviewers for their constructive comments.

This work³ was funded by Vinnova and Wallenberg AI, Autonomous Systems and Software Program (WASP), Sweden (research grant). The computations were enabled by the supercomputing resource Berzelius provided by National Supercomputer Centre at Linköping University and the Knut and Alice Wallenberg foundation, Sweden.

References

- [1] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *International Journal of Robotics Research (IJRR)*, 2013. 1
- [2] Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2446–2454, 2020. 1
- [3] Tianyu Liu, Qing hai Liao, Lu Gan, Fulong Ma, Jie Cheng, Xupeng Xie, Zhe Wang, Yingbing Chen, Yilong Zhu, Shuyang Zhang, et al. The role of the hercules autonomous vehicle during the covid-19 pandemic: An autonomous logistic vehicle for contactless goods transportation. *IEEE Robotics & Automation Magazine*, 28(1):48–58, 2021. 1
- [4] Dian Chen and Philipp Krähenbühl. Learning from all vehicles. In *CVPR*, 2022. 1
- [5] Ionel Gog, Sukrit Kalra, Peter Schafhalter, Matthew A Wright, Joseph E Gonzalez, and Ion Stoica. Pylot: A modular platform for exploring latency-accuracy tradeoffs in autonomous vehicles. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8806–8813. IEEE, 2021. 1
- [6] Qingwen Zhang, Mingkai Tang, Ruoyu Geng, Feiyi Chen, Ren Xin, and Lujia Wang. Mmfnet: Multi-modal-fusionnet for end-to-end driving. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8638–8643. IEEE, 2022. 1
- [7] Xiaosong Jia, Penghao Wu, Li Chen, Jiangwei Xie, Conghui He, Junchi Yan, and Hongyang Li. Think twice before driving: Towards scalable decoders for end-to-end autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21983–21994, 2023. 1
- [8] Yihan Hu, Jiazhi Yang, Li Chen, Keyu Li, Chonghao Sima, Xizhou Zhu, Siqi Chai, Senyao Du, Tianwei Lin, Wenhai Wang, et al. Planning-oriented autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17853–17862, 2023. 1, 2
- [9] Xiaosong Jia, Yulu Gao, Li Chen, Junchi Yan, Patrick Langechuan Liu, and Hongyang Li. Driveadapter: Breaking the coupling barrier of perception and planning in end-to-end autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7953–7963, 2023. 1
- [10] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Yu Qiao, and Jifeng Dai. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. In *European conference on computer vision*, pages 1–18. Springer, 2022. 1
- [11] Long Chen, Yuchen Li, Chao Huang, Bai Li, Yang Xing, Daxin Tian, Li Li, Zhongxu Hu, Xiaoxiang Na, Zixuan Li, et al. Milestones in autonomous driving and intelligent vehicles: Survey of surveys. *IEEE Transactions on Intelligent Vehicles*, 8(2):1046–1056, 2022. 1
- [12] Can Cui, Yunsheng Ma, Xu Cao, Wenqian Ye, and Ziran Wang. Drive as you speak: Enabling human-like interaction with large language models in autonomous vehicles. *arXiv preprint arXiv:2309.10228*, 2023. 1, 3
- [13] Can Cui, Yunsheng Ma, Xu Cao, Wenqian Ye, and Ziran Wang. Receive, reason, and react: Drive as you say with large language models in autonomous vehicles. *arXiv preprint arXiv:2310.08034*, 2023. 1, 3
- [14] Wayve Research. Lingo-1: Exploring natural language for autonomous driving. <https://wayve.ai/thinking/lingo-natural-language-autonomous-driving>, 2023. 1
- [15] DriveLM Contributors. Drivelm: Drive on language. <https://github.com/OpenDriveLab/DriveLM>, 2023. 1
- [16] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*, 2022. 1
- [17] Nan-ning Zheng, Zi-yi Liu, Peng-ju Ren, Yong-qiang Ma, Shi-tao Chen, Si-yu Yu, Jian-ru Xue, Ba-dong Chen, and Fei-yue Wang. Hybrid-augmented intelligence: collaboration and cognition. *Frontiers of Information Technology & Electronic Engineering*, 18(2):153–179, 2017. 1
- [18] 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, 3
- [19] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021. 1

³We have used ChatGPT for editing and polishing author-written text.

- [20] 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, 3
- [21] OpenAI. Gpt-4 technical report, 2023. 1, 2, 3
- [22] Andy Zeng, Maria Attarian, Krzysztof Marcin Choromanski, Adrian Wong, Stefan Welker, Federico Tombari, Aavek Purohit, Michael S Ryoo, Vikas Sindhwani, Johnny Lee, et al. Socratic models: Composing zero-shot multimodal reasoning with language. In *The Eleventh International Conference on Learning Representations*, 2022. 1
- [23] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837, 2022. 1, 2
- [24] Namgyu Ho, Laura Schmid, and Se-Young Yun. Large language models are reasoning teachers. *arXiv preprint arXiv:2212.10071*, 2022. 1
- [25] Maria Tsimpoukelli, Jacob L Menick, Serkan Cabi, SM Eslami, Oriol Vinyals, and Felix Hill. Multimodal few-shot learning with frozen language models. *Advances in Neural Information Processing Systems*, 34:200–212, 2021. 1
- [26] Ziyi Zhang, Weikai Chen, Hui Cheng, Zhen Li, Siyuan Li, Liang Lin, and Guanbin Li. Divide and contrast: Source-free domain adaptation via adaptive contrastive learning. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 5137–5149. Curran Associates, Inc., 2022. 1
- [27] Efstratios Gavves, Thomas Mensink, Tatiana Tommasi, Cees GM Snoek, and Tinne Tuytelaars. Active transfer learning with zero-shot priors: Reusing past datasets for future tasks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2731–2739, 2015. 1
- [28] Devin Schwab, Yifeng Zhu, and Manuela Veloso. Zero shot transfer learning for robot soccer. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems*, pages 2070–2072, 2018. 1
- [29] Wei Ying, Yu Zhang, Junzhou Huang, and Qiang Yang. Transfer learning via learning to transfer. In *International Conference on Machine Learning*, pages 5085–5094. PMLR, 2018. 1
- [30] Jae Woong Soh, Sunwoo Cho, and Nam Ik Cho. Meta-transfer learning for zero-shot super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3516–3525, 2020. 1
- [31] Sreyasi Nag Chowdhury, Ruwan Wickramarachchi, Mohamed Hassan Gad-Elrab, Daria Stepanova, and Cory Henson. Towards leveraging commonsense knowledge for autonomous driving. In *20th International Semantic Web Conference*, pages 1–5. CEUR-WS. org, 2021. 2, 3
- [32] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021. 2
- [33] Dongming Wu, Wencheng Han, Tiancai Wang, Yingfei Liu, Xiangyu Zhang, and Jianbing Shen. Language prompt for autonomous driving. *arXiv preprint arXiv:2309.04379*, 2023. 2
- [34] 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
- [35] 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*, pages 1043–1052, 2023. 2
- [36] Bu Jin, Xinyu Liu, Yupeng Zheng, Pengfei Li, Hao Zhao, Tong Zhang, Yuhang Zheng, Guyue Zhou, and Jingjing Liu. Adapt: Action-aware driving caption transformer. *arXiv preprint arXiv:2302.00673*, 2023. 2
- [37] Bojan Mrazovac and Milan Z Bjelica. Human-centric role in self-driving vehicles: Can human driving perception change the flavor of safety features? *IEEE Intelligent Transportation Systems Magazine*, 15(1):117–125, 2022. 2
- [38] Guofa Li, Liang Yang, Shen Li, Xiao Luo, Xingda Qu, and Paul Green. Human-like decision making of artificial drivers in intelligent transportation systems: An end-to-end driving behavior prediction approach. *IEEE Intelligent Transportation Systems Magazine*, 14(6):188–205, 2021. 2
- [39] 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. 2
- [40] Xingcheng Zhou, Mingyu Liu, Bare Luka Zagar, Ekim Yurtsever, and Alois C Knoll. Vision language models in autonomous driving and intelligent transportation systems. *arXiv preprint arXiv:2310.14414*, 2023. 3
- [41] Xinpeng Ding, Jianhua Han, Hang Xu, Wei Zhang, and Xiaomeng Li. Hilm-d: Towards high-resolution understanding in multimodal large language models for autonomous driving. *arXiv preprint arXiv:2309.05186*, 2023. 3
- [42] Kanishk Jain, Varun Chhangani, Amogh Tiwari, K Madhava Krishna, and Vineet Gandhi. Ground then navigate: Language-guided navigation in dynamic scenes. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4113–4120. IEEE, 2023. 3
- [43] Xuejiao Zhao. Explainable q&a system based on domain-specific knowledge graph. 2021. 3
- [44] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov,

Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023. 3

[45] Chao Zheng, Kun Tang, Zhipeng Cao, Xu Cao, Yunsheng Ma, Can Cui, Wenqian Ye, Ziran Wang, Shawn Mei, and Tong Zhou. Challenge: In-cabin user command understanding (ucu). <https://github.com/llvm-ad/ucu-dataset>, 2023. 3

[46] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 3