

ADAPT: Action-aware Driving Caption Transformer

Bu Jin^{1,2}, Xinyu Liu^{2,4}, Yupeng Zheng^{1,2}, Pengfei Li³, Hao Zhao²✉, Tong Zhang⁵,
Yuhang Zheng^{2,6}, Guyue Zhou² and Jingjing Liu²

Abstract—End-to-end autonomous driving has great potential in the transportation industry. However, the lack of transparency and interpretability of the automatic decision-making process hinders its industrial adoption in practice. There have been some early attempts to use attention maps or cost volume for better model explainability which is difficult for ordinary passengers to understand. To bridge the gap, we propose an end-to-end transformer-based architecture, ADAPT (Action-aware Driving cAPtion Transformer), which provides *user-friendly* natural language narrations and reasoning for each decision making step of autonomous vehicular control and action. ADAPT jointly trains both the driving caption task and the vehicular control prediction task, through a shared video representation. Experiments on BDD-X (Berkeley DeepDrive eXplanation) dataset demonstrate state-of-the-art performance of the ADAPT framework on both automatic metrics and human evaluation. To illustrate the feasibility of the proposed framework in real-world applications, we build a novel deployable system that takes raw car videos as input and outputs the action narrations and reasoning in real time. The code, models and data are available at <https://github.com/jxbbb/ADAPT>.

I. INTRODUCTION

The goal of an autonomous system is to gain precise perception of the environment, make safe real-time decisions, take reliable actions without human involvement and provide a safe and comfortable ride experience for passengers. There are generally two types of paradigms for autopilot controller design: mediation-aware method [1], [2] and end-to-end learning approach [3]–[23]. Mediation-aware approaches rely on recognizing human-specified features such as vehicles, lane markings, etc., which require rigorous parameter tuning to achieve satisfactory performance. In contrast, end-to-end methods directly take raw data from sensors as input to generate planning routes or control signals.

One of the key challenges in deploying such autonomous control systems to real vehicles is that intelligent decision-making policies in autonomous cars are often too complicated and difficult for common passengers to understand, for whom the safety of such vehicles and their controllability is the top priority.

Some previous work has explored the interpretation of autonomous navigation [13], [14], [24]–[30]. Cost map, for

example, is employed in [13] to interpret the actions of a self-driving system by visualizing the difficulty of traversing through different areas of the map. Visual attention is utilized in [24] to filter out non-salient image regions, and [31] constructs BEV (Bird’s eye view) to visualize the motion information of the vehicle. However, these interfaces can easily lead to misinterpretation if the user is unfamiliar with the system.

An ideal solution is to include natural language narrations to guide the use throughout the decision making and action taking process of the autonomous control module, which is comprehensible and user-friendly. Furthermore, an additional reasoning explanation for each control/action decision can help users understand the current state of the vehicle and the surrounding environment, as supporting evidence for the actions taken by the autonomous vehicle. For example, “[*Action narration*:] the car pulls over to the right side of the road, [*Reasoning*:] because the car is parking”, as shown in Fig. 1. Explaining vehicle behaviors via natural language narrations and reasoning thus makes the whole autonomous system more transparent and easier to understand.

To this end, we propose ADAPT, the first action-aware transformer-based driving action captioning architecture that provides for passengers user-friendly natural language narrations and reasoning of autonomous driving vehicles. To eliminate the discrepancy between the captioning task and the vehicular control signal prediction task, we jointly train these two tasks with a shared video representation. This multi-task framework can be built upon various end-to-end autonomous systems by incorporating a text generation head.

We demonstrate the effectiveness of the ADAPT approach on a large-scale dataset that consists of control signals and videos along with action narration and reasoning. Based on ADAPT, we build a novel deployable system that takes raw vehicular navigation videos as input and generates the action narrations and reasoning explanations in real time.

Our contributions can be summarized as:

- We propose ADAPT, a new end-to-end transformer-based action narration and reasoning framework for self-driving vehicles.
- We propose a multi-task joint training framework that aligns both the driving action captioning task and the control signal prediction task.
- We develop a deployable pipeline for the application of ADAPT in both the simulator environment and the real world.

¹Institute of Automation, Chinese Academy of Sciences, China, jinbu2022@ia.ac.cn.

²Institute for AI Industry Research (AIR), Tsinghua University, China, zhaohao@air.tsinghua.edu.cn.

³Department of Computer Science and Technology, Tsinghua University, China

⁴Xidian University, China

⁵Southern University of Science and Technology, China

⁶School of Mechanical Engineering and Automation, Beihang University, China

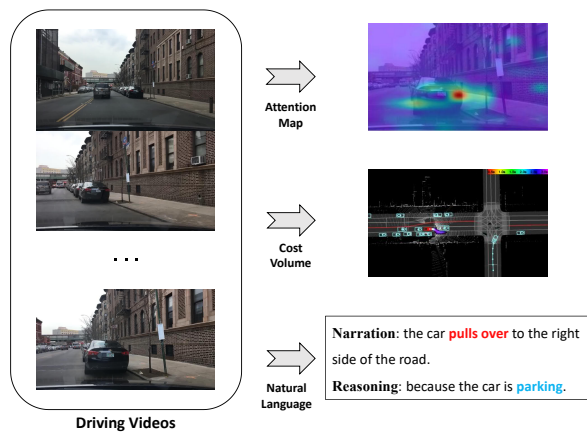


Fig. 1. Different interpretation methods of self-driving vehicles, including attention map [24], cost volume [13] and natural language. Despite the effectiveness of attention map or cost volume, language-based interpretation is more user-friendly to common passengers.

II. RELATED WORK

A. Video Captioning

The main goal of the video captioning task is to describe the objects and their relationship of a given video in natural language. Early researches [32]–[35] generate sentences with specific syntactic structures by filling recognized elements in fixed templates, which are inflexible and lack of richness. [36]–[45] exploit sequence learning approaches to generate natural sentences with flexible syntactic structures. Specifically, these methods employ a video encoder to extract frame features and a language decoder to learn visual-textual alignment for caption generation. To introduce detailed objects and actions to caption, [46]–[48] exploit object-level representations that capture detailed object-aware interaction features in videos. [49] further develops a novel dual-branch convolutional encoder to jointly learn the content and semantic information of videos. Moreover, [50] adapts the unimodal transformer to video captioning and employs a sparse boundary-aware pooling to reduce the redundancy in video frames. The development of scene understanding [51]–[62] also contribute a lot to the captioning task. Most recently, [63] proposes an end-to-end transformer-based model SWIN-BERT, which utilizes a sparse attention mask to lessen the redundant and irrelevant information in consecutive video frames.

While existing architectures achieve promising results for general video captioning, it cannot be directly applied to action representation because simply transferring video caption to self-driving action representation would miss some key information like the speed of the vehicle, which is essential in the autonomous system. How to effectively use these multi-modal information to generate sentences remains a mystery, which is the focus of our work.

B. End-to-End Autonomous Driving

Learning-based autonomous driving is an active research area [64], [65]. Some learning-based driving methods such

as affordances [3], [4] and reinforcement learning [5]–[7] are employed, gaining promising performance. Imitation methods [8]–[13] are also utilized to regress the control commands from human demonstrations. For example, [14]–[16] model the future behavior of driving agents like vehicles, cyclists or pedestrians to predict the vehicular waypoints, while [17]–[23] predict vehicular control signals directly according to the sensor input, which is similar to our control signal prediction sub-task.

C. Interpretability of Autonomous Driving

Interpretability, or the ability to provide a comprehensive explanation plays a significant role in the social acceptance of artificial intelligence [66], [67] and autonomous driving is no exception. Most interpretable approaches of autonomous vehicles are vision-based [24]–[27], [31] or LiDAR-based [13], [14], [28]–[30]. [24] first utilizes the visualization of an attention map that filters out non-salient image regions to make autonomous vehicles reasonable and interpretable. Nevertheless, the attention map may easily include some less important areas which cause misunderstanding for passengers. [25]–[27], [31] constructs BEV (Bird’s eye view) from a vehicle camera to visualize the motion information and environmental status of the vehicle. [13] takes as input LiDAR and HD maps to forecast the bounding boxes of driving agents and exploits cost volume to explain the reason for the planner’s decision. Furthermore, [14] constructs an online map from segmentation as well as the states of driving agents to avoid heavy dependence on HD maps.

Although the vision-based or LiDAR-based approaches provide promising results, the lack of linguistic interpretation makes them too complicated for passengers like the elderly to understand. [68] first explores the possibility of textual explanations for self-driving vehicles, which offline extracts video features from control signal prediction task and conducts video captioning afterwards. Unfortunately, the discrepancy between these two tasks makes the offline-extracted features sub-optimal for downstream captioning task, which is the focus of our work.

D. Multi-task Learning in Autonomous Driving

ADAPT leverage the concept of multi-task learning, in which the model is trained on a joint objective of text generation and control signal prediction. Multi-task learning helps extract more useful information by exploiting inductive biases between different tasks [69] and has shown promising prospects in autonomous driving. [70], [71] shows that detection and tracking can be trained together. [72] further applies a joint detector and trajectory predictor into a single model and gains promising results. This idea is extended by [73] to simultaneously predict the intention of actors. More recently, [13] further includes a cost map based control signal planner in the joint model. These works show that joint training of different tasks improves the performance of individual tasks due to better data utilization and shared features, which inspires our joint training strategy of control signal prediction task and text generation task.

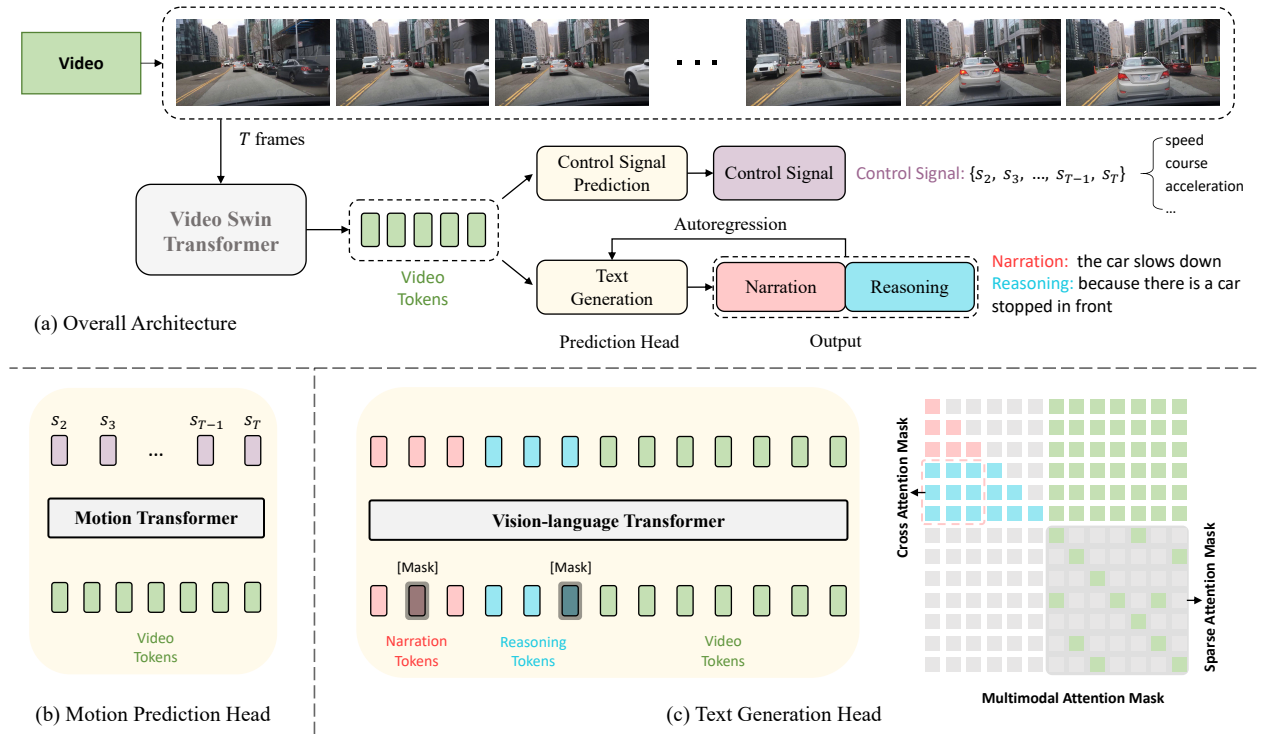


Fig. 2. Overview of ADAPT framework. (a) Input is a vehicle-front-view video, and outputs are predicted vehicle’s control signals and the narration and reasoning of the current action. We first densely and uniformly sample T frames from the video, which are sent to the learnable video swin transformer and tokenized into video tokens. Different prediction heads generate final motion results and text results. (b)(c) show the prediction heads, respectively.

III. METHOD

A. Overview

The ADAPT architecture is illustrated in Fig. 2, which addresses two tasks: Driving Caption Generation (DCG) and Control Signal Prediction (CSP). DCG takes a sequence of raw video frames as inputs, and outputs two natural language sentences: one describes the vehicle’s action (e.g., “the car is accelerating”), and the other explains the reasoning for taking this action (e.g., “because the traffic lights turn green”). CSP takes the same video frames as inputs, and outputs a sequence of control signals, such as speed or acceleration.

Generally, DCG and CSP tasks share the same Video Encoder, while employing different heads to generate the corresponding results. For DCG task, we employ a vision-language transformer encoder to produce two natural language sentences via sequence-to-sequence generation. For CSP task, we use a motion transformer encoder to predict the control signal sequence.

B. Video Encoder

Following Swinbert [63], we employ Video Swin Transformer (video swin) [74] as the visual encoder, which encodes the frames from a vehicular video into video features. More specifically, given a vehicular video captured from the first perspective, we first do uniform sampling to get T frames of size $H \times W \times 3$. These frames are passed as inputs to video swin, resulting in feature F_V of size

$\frac{T}{2} \times \frac{H}{32} \times \frac{W}{32} \times 8C$, where C is the channel dimension defined in video swin. Then the video features are fed into different prediction heads for individual tasks.

C. Prediction Heads

Text Generation Head The purpose of the text generation head is to generate two sentences that describe both the action of the vehicle and the reason behind it. As mentioned in Sec. III-B, the video frames are encoded to video features F_V of size $\frac{T}{2} \times \frac{H}{32} \times \frac{W}{32} \times 8C$. Then we create the video tokens by flattening the video features except channel dimension and tokenizing along that dimension, which results in $\frac{T}{2} \times \frac{H}{32} \times \frac{W}{32}$ tokens with dimension of $8C$. As for the text inputs (action narrations and reasoning), we first tokenize each sentence and pad it to a fixed length. Then we concatenate these two sentences and embed them with an embedding layer. To identify the difference between action narration and reasoning, we exploit a segment embedding method (widely used in Bert [75]) to distinguish them. And we use a learnable MLP that transforms the dimension of video tokens to ensure the dimension consistency between video tokens and text tokens. Finally, the text tokens and video tokens are fed into the vision-language transformer encoder, which will generate a new sequence includes both action narrations and reasoning.

Control Signal Prediction Head The goal of CSP head is to predict the control signals (e.g. acceleration) of the vehicle based on video frames. Given video features of T frames, along with the corresponding control signal recordings $\mathcal{S} = \{s_1, s_2, \dots, s_T\}$, the output of CSP head is a sequence of

control signals $\hat{S} = \{\hat{s}_2, \dots, \hat{s}_T\}$. Each control signal s_i or \hat{s}_i is a n -tuple, where n refers to how many types of sensor we exploit. We first tokenize the video features, then utilize another transformer (motion transformer) to generate the prediction of these control signals. The loss function \mathcal{L}_{CSP} is defined as the mean squared error of S and \hat{S} :

$$\mathcal{L}_{CSP} = \frac{1}{T-1} \sum_{i=2}^T (s_i - \hat{s}_i)^2 \quad (1)$$

Note that we do not predict control signal corresponding to the first frame, since the dynamic information of the first frame is limited, while other signals can be easily inferred from previous frames.

D. Joint Training

In our framework, we assume that CSP and DCG tasks are aligned on the semantic level of the video representation. Intuitively, action narration and the control signal data are different expression forms of the action of self-driving vehicles, while reasoning explanation concentrates on the elements of the environment that influence the action of the vehicles. We believe that jointly training these tasks in a single network can improve performance by leveraging the inductive biases between different tasks.

During training, CSP and DCG are performed jointly. We simply add the \mathcal{L}_{CSP} and \mathcal{L}_{DCG} to get the final loss function:

$$\mathcal{L} = \mathcal{L}_{CSP} + \mathcal{L}_{DCG} \quad (2)$$

Despite the joint training of both tasks, inference on each task can be carried out independently. For the DCG task, ADAPT takes a video sequence as input, and outputs the driving caption with two segments. Text generation is performed in an auto-regressive manner. Specifically, our model starts with a "[CLS]". In each recursion, it uses previously generated words as the inputs and generates next word. Generation continues until the model predicts the ending word "[SEP]" or the sentence exceeds the permitted length. After padding the first sentence to the maximum length, we concatenate another "[CLS]" to the inputs and repeat the aforementioned process, then two separate sentences are generated.

IV. EXPERIMENT

In this section, we evaluate ADAPT over metrics of the standard captioning task, including BLEU4 [76], METEOR [77], ROUGE-L [78] and CIDEr [79] (abbreviated as B4, M, R and C in later tables). As quantitative evaluation of captioning is still an open question, we also provide detailed human evaluation results for the subjective correctness of the generated text. Ablation studies further demonstrate the effectiveness of the proposed joint-training framework.

A. Dataset

BDD-X [68] is a driving-domain caption dataset, consisting of nearly 7000 videos paired with control signals. The videos and control signals are collected from BDD100K

TABLE I
COMPARISON WITH SOTA ON VIDEO CAPTIONING METRICS

Method	Narration			Reasoning		
	B4	C	M	B4	C	M
S2VT [42]	30.2	179.8	27.5	6.3	53.4	11.2
S2VT++ [42]	27.1	157.0	26.4	5.8	52.7	10.9
SAA [68]	31.8	214.8	29.1	7.1	66.1	12.2
WAA [68]	32.3	215.8	29.2	7.3	69.5	12.2
Ours	34.6	247.5	30.6	11.4	102.6	15.2

TABLE II
COMPARISON WITH SOTA ON HUMAN EVALUATION

Method	Narration	Reasoning	Full sentence
SAA [68]	90.8%	62.4%	-
WAA [68]	93.5%	66.0%	-
Ours	90.0%	90.3%	82.7%

dataset [80]. Each video has a duration of 40 seconds on average, with 1280×720 resolution and 30 FPS. Each video contains 1 to 5 vehicle behaviors, such as accelerating, turning right or merging lanes. All these behaviors are accompanied by text annotation, including action narration (e.g., "the car stops") and reasoning (e.g., "because the traffic light is red"). There are around 29000 behavior-annotation pairs in total. To the best of our knowledge, BDD-X is the only driving-domain caption dataset accompanied by car videos and control signals.

B. Implementation Details

The video swin transformer is pre-trained on Kinetics-600 [81], while the vision-language transformer and motion transformer are randomly initialized. Note that in our implementation we do not freeze the parameters of video swin, so ADAPT is trained in a complete end-to-end manner. The input video frames are resized and cropped to the spatial size of 224. And for narration and reasoning, we use WordPiece embeddings [75] instead of the whole words (e.g., "stops" is cut to "stop" and "#s") and the maximal length of each sentence is 15. During training period, we randomly mask 50% of the tokens for masked language modeling. And the masked token has 80% chance to be a "[MASK]" token, 10% chance to be a random word, and 10% chance to remain the same. We utilize AdamW optimizer and adopt a warm-up learning rate during the training process. The whole training for 40 epochs takes about 13 hours on 4 NVIDIA V100 GPUs with a batch size of 4 per GPU.

C. Main Results

We compare ADAPT with state-of-the-art methods on BDD-X dataset. Table I shows the comparison results on standard captioning metrics. We observe that ADAPT achieves significant performance gain over existing methods. Specifically, ADAPT outperforms prior state-of-the-art work [68] by 31.7 for action narration and 33.1 for reasoning on CIDEr metric.

In addition to automatic evaluation measures, we also conduct human evaluation to measure the subjective correctness of output narration and reasoning. The whole evaluation

TABLE III
SINGLE CAPTIONING VS. ACTION-AWARE CAPTIONING

Method	Narration				Reasoning			
	B4	C	M	R	B4	C	M	R
Single	33.2	238.9	29.7	62.0	8.6	89.7	14.1	31.4
Single+	33.9	248.3	30.5	63.1	9.3	97.2	14.6	31.5
Ours	34.6	247.5	30.6	62.8	11.4	102.6	15.2	32.0

TABLE IV
ANALYSIS ON CONTROL SIGNAL TYPES

Speed	Signals	Narration			Reasoning		
	Course	C	M	R	C	M	R
✓		232.0	29.9	61.5	88.0	15.1	31.0
	✓	218.2	29.3	61.2	88.6	14.1	30.6
✓	✓	247.5	30.6	62.8	102.6	15.2	32.0

process is divided into three sections: (1) narration, (2) reasoning, and (3) full sentence. During the first section, a human evaluator judges whether the predicted narrations conform to the vehicle’s action. In the second section, we display both ground-truth narration and predicted reasoning, and require human evaluators to judge whether the reasoning is correct. Then in the last section, both predicted narrations and predicted reasoning are displayed. Table II shows that ADAPT outperforms previous work in reasoning accuracy while maintaining high accuracy on narration evaluation, demonstrating the effectiveness of ADAPT.

D. Ablation Study

We conduct a comprehensive ablation study to analyze various aspects of ADAPT design.

Effect of Action-aware Joint Training To investigate the effect of action-awareness in joint training on ADAPT, we train a single captioning model by removing the CSP (control signal prediction) head of ADAPT, referred to as "Single". As shown in Table III, ADAPT outperforms single training with an improvement of 15.9 for narration and 7.2 for reasoning on CIDEr metric. This suggests that cues from the other task help regularize the shared video representation and improve the performance of the text generation task.

Additionally, we can see from Fig. 2(a) that the caption and control signal data are employed in two streams in ADAPT. An interesting question is: can we simply pass the control signals to the multi-modal transformer to get the final caption prediction? So we create such an architecture that

TABLE V
IMPACT OF INTERACTION BETWEEN NARRATION AND REASONING

Method	Narration			Reasoning		
	B4	C	M	B4	C	M
w/o cross attn	32.7	234.8	30.0	10.7	96.6	15.1
w/ swapped attn	28.3	180.4	28.7	9.3	97.7	14.3
Ours	34.6	247.5	30.6	11.4	102.6	15.2
Narration only	32.9	240.4	29.6	-	-	-
Reasoning only	-	-	-	8.1	94.4	13.3

TABLE VI
IMPACT OF VIDEO FRAMES(T)

Method	Narration				Reasoning				Cost(min)
	B4	C	M	R	B4	C	M	R	
2	33.4	227.7	28.7	61.0	8.7	62.9	15.1	29.8	294
4	32.9	225.7	29.0	60.9	9.9	81.3	14.9	31.1	382
8	32.6	236.1	29.3	61.8	8.4	83.7	13.4	30.6	447
16	32.5	231.0	29.5	61.9	8.7	91.5	13.8	32.0	528
32	34.6	247.5	30.6	62.8	11.4	102.6	15.2	32.0	797

takes video tokens, control signal tokens (generated by a learnable embedding layer) and masked text tokens as input and generates predictions of the masked tokens, which is referred to as "Single+". Results are shown in the second row of Table III. We can see that the proposed ADAPT still achieves the best results, especially for reasoning segment, which demonstrates the superiority of multi-task learning over using both videos and control signals as inputs despite the latter is an intuitive setting.

Impact of Different Control Signal Types In our implementations, we leverage control signals (e.g., course) as supervision for the CSP task. In this analysis, we investigate the impact of different supervision signal types of ADAPT. The base signals in our experiments are speed and course. We first conduct experiments by removing one of them, results of which are shown in the first two rows of Table IV. Then in the third row both speed and course are utilized, which is the same as previous experiments. We observe that the removal of each signal leads to the decrease of performance. For example, the CIDEr metric decreases by 29.3 for narration and by 14.0 for reasoning without the speed inputs. This is understandable because being aware of speed and course can help the network learn representations that are informative for narration and reasoning and the lack of either can result in the bias of video representations.

Interaction between Narration and Reasoning Compared with the general caption task, the driving caption task generates two sentences: action narration and reasoning. In this section, we explore how these two segments interact with each other by controlling the attention mask or the order of two sentences.

Specifically, as shown in the right of Fig. 2(c), we use a causal self-attention mask for each sentence where a word token can only attend to the existing output tokens, and employ sparse attention [63] for video tokens. The reasoning segment has full attention to the narration segment, referred to as cross attention, which defines the dependence of reasoning on narration. In this section, we first conduct experiments without cross attention or with swapped cross attention (by swapping the order of narration and reasoning). Results are reported in Table V. Compared with the default setting (denoted as "Ours"), results without cross attention have lower performance in both sentences, which indicates that conditioning the reasoning segment on the narration segment is beneficial for training. And the performance with swapped cross attention also decreases, especially for the narration part, which further demonstrates this dependence of reasoning on narration, instead of the other way around.

TABLE VII
COMPARISON ON CONTROL SIGNALS PREDICTION ACCURACY

Method	Course					Speed						
	RMSE(degree)↓	$A_{0.1}$ ↑	$A_{0.5}$ ↑	$A_{1.0}$ ↑	$A_{5.0}$ ↑	$A_{10.0}$ ↑	RMSE(m/s)↓	$A_{0.1}$ ↑	$A_{0.5}$ ↑	$A_{1.0}$ ↑	$A_{5.0}$ ↑	$A_{10.0}$ ↑
Single	6.3	8.3	84.7	90.5	97.2	98.7	3.4	5.0	25.5	37.8	86.8	98.7
Ours	6.4	62.2	85.5	89.9	97.2	98.8	2.5	11.1	28.1	45.3	94.3	99.5

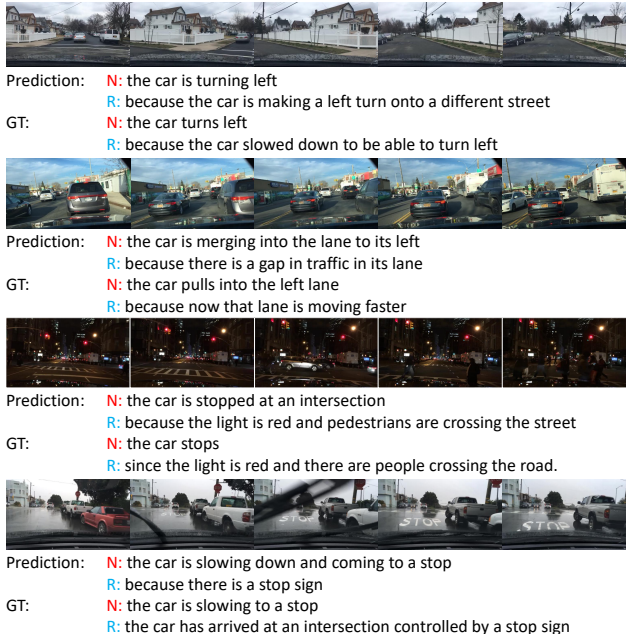


Fig. 3. Qualitative analysis: generated narrations correctly describe the current action, with sound reasoning.

Additionally, we conduct experiments with only one sentence, referred to as "Narration only" and "Reasoning only". Table V shows that training with both sentences yields improvement on the performance, especially for the reasoning segment, indicating that the interaction between narration and reasoning promotes each component of the full caption task.

Impact of Different Sampling Rates In previous experiments, we uniformly sample $T = 32$ frames from a given video, along with control signal data of the same timestamp. In this study, we investigate the impact of sampling rate by varying the number of sampled frames. Specifically, we uniformly sample $T = 2, 4, 8, 16, 32$ frames from a variable-length video, as shown in Table VI. The performance of ADAPT improves steadily as the sampled number increases since more frames lead to less missing visual content. This suggests that caption results can be enhanced by densely sampled frames and control signals. The training time costs are also provided in Table VI. We hope this ablation provides robotics practitioner with insights about the accuracy-efficiency trade-off of driving caption.

E. Analysis on Control Signal Prediction

Although the main goal of driving caption task is to generate sentences, we also investigate the performance of control signal prediction tasks. We employ root mean squared

error (RMSE) and a tolerant accuracy (A_σ) to measure the final performance. Tolerant accuracy means we first use two thresholds to determine the range of the control signal deviation and truncate it. For example, we define the truncation value of predicted course \hat{c} as:

$$c_\sigma = \begin{cases} 1, & -\sigma < \hat{c} - c < \sigma \\ 0, & otherwise \end{cases} \quad (3)$$

where c is the ground-truth course and σ is the tolerant threshold value. Then A_σ of course represents the accuracy of c_σ recorded as a percentage, and A_σ of speed is defined similarly. Results are provided in Table VII. We observe that our joint training framework can further improve the performance of control signal prediction, indicating the benefit of joint training.

F. Deployment in Autonomous Systems

We further develop a pipeline for the deployment of ADAPT in both the simulator environment (e.g., Carla [82]) and the real world. The system takes raw vehicular videos as input and generates action narrations and reasoning explanations in real time. Specifically, we first record the frames captured by the camera from the front view. Then the frames in the last several seconds are passed as input to ADAPT to generate the action narration and reasoning of the current step. Moreover, we further utilize text-to-speech technology to convert the generated sentences into speech narration, to make it more convenient and more interactive for common passengers (especially helpful for vision-impaired passengers).

V. CONCLUSION

Language-based interpretability is essential for the social acceptance of self-driving vehicles. We present Adapt (Action-aware Driving cAPtion Transformer), a new end-to-end transformer-based framework for generating action narration and reasoning for self-driving vehicles. ADAPT utilizes multi-task joint training to reduce the discrepancy between the driving action captioning task and the control signal prediction task. Experiments on BDD-X dataset over standard captioning metrics as well as human evaluation demonstrate the effectiveness of ADAPT over state-of-the-art methods. We further develop a deployable pipeline for the application of ADAPT in both simulator environment and the real world.

VI. ACKNOWLEDGEMENT

This work was supported by Baidu Inc. through Apollo-AIR Joint Research Center.

REFERENCES

- [1] S. Behere and M. Törngren, "A functional architecture for autonomous driving," in *Proceedings of the First International Workshop on Automotive Software Architecture*, 2015, pp. 3–10.
- [2] E. Yurtsever, J. Lambert, A. Carballo, and K. Takeda, "A survey of autonomous driving: Common practices and emerging technologies," *IEEE access*, vol. 8, pp. 58 443–58 469, 2020.
- [3] A. Sauer, N. Savinov, and A. Geiger, "Conditional affordance learning for driving in urban environments," in *Conference on Robot Learning*. PMLR, 2018, pp. 237–252.
- [4] Y. Xiao, F. Codevilla, C. Pal, and A. M. López, "Action-based representation learning for autonomous driving," *arXiv preprint arXiv:2008.09417*, 2020.
- [5] D. Chen, V. Koltun, and P. Krähenbühl, "Learning to drive from a world on rails," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 15 590–15 599.
- [6] M. Toromanoff, E. Wirbel, and F. Moutarde, "End-to-end model-free reinforcement learning for urban driving using implicit affordances," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 7153–7162.
- [7] J. Wang, Y. Wang, D. Zhang, Y. Yang, and R. Xiong, "Learning hierarchical behavior and motion planning for autonomous driving," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 2235–2242.
- [8] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang *et al.*, "End to end learning for self-driving cars," *arXiv preprint arXiv:1604.07316*, 2016.
- [9] F. Codevilla, M. Müller, A. López, V. Koltun, and A. Dosovitskiy, "End-to-end driving via conditional imitation learning," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 4693–4700.
- [10] M. Müller, A. Dosovitskiy, B. Ghanem, and V. Koltun, "Driving policy transfer via modularity and abstraction," *arXiv preprint arXiv:1804.09364*, 2018.
- [11] D. A. Pomerleau, "Alvinn: An autonomous land vehicle in a neural network," *Advances in neural information processing systems*, vol. 1, 1988.
- [12] B. Wei, M. Ren, W. Zeng, M. Liang, B. Yang, and R. Urtasun, "Perceive, attend, and drive: Learning spatial attention for safe self-driving," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 4875–4881.
- [13] W. Zeng, W. Luo, S. Suo, A. Sadat, B. Yang, S. Casas, and R. Urtasun, "End-to-end interpretable neural motion planner," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 8660–8669.
- [14] S. Casas, A. Sadat, and R. Urtasun, "Mp3: A unified model to map, perceive, predict and plan," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 14 403–14 412.
- [15] D. Chen, B. Zhou, V. Koltun, and P. Krähenbühl, "Learning by cheating," in *Conference on Robot Learning*. PMLR, 2020, pp. 66–75.
- [16] A. Filos, P. Tigkas, R. McAllister, N. Rhinehart, S. Levine, and Y. Gal, "Can autonomous vehicles identify, recover from, and adapt to distribution shifts?" in *International Conference on Machine Learning*. PMLR, 2020, pp. 3145–3153.
- [17] A. Behl, K. Chitta, A. Prakash, E. Ohn-Bar, and A. Geiger, "Label efficient visual abstractions for autonomous driving," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 2338–2345.
- [18] A. Bühler, A. Gaidon, A. Cramariuc, R. Ambrus, G. Rosman, and W. Burgard, "Driving through ghosts: Behavioral cloning with false positives," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 5431–5437.
- [19] F. Codevilla, E. Santana, A. M. López, and A. Gaidon, "Exploring the limitations of behavior cloning for autonomous driving," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 9329–9338.
- [20] Z. Huang, C. Lv, Y. Xing, and J. Wu, "Multi-modal sensor fusion-based deep neural network for end-to-end autonomous driving with scene understanding," *IEEE Sensors Journal*, vol. 21, no. 10, pp. 11 781–11 790, 2020.
- [21] E. Ohn-Bar, A. Prakash, A. Behl, K. Chitta, and A. Geiger, "Learning situational driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11 296–11 305.
- [22] A. Prakash, A. Behl, E. Ohn-Bar, K. Chitta, and A. Geiger, "Exploring data aggregation in policy learning for vision-based urban autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11 763–11 773.
- [23] Y. Xiao, F. Codevilla, A. Gurram, O. Urfalioglu, and A. M. López, "Multimodal end-to-end autonomous driving," *IEEE Transactions on Intelligent Transportation Systems*, 2020.
- [24] J. Kim and J. Canny, "Interpretable learning for self-driving cars by visualizing causal attention," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2942–2950.
- [25] A. Hu, Z. Murez, N. Mohan, S. Dudas, J. Hawke, V. Badrinarayanan, R. Cipolla, and A. Kendall, "Fiery: Future instance prediction in bird's-eye view from surround monocular cameras," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 15 273–15 282.
- [26] J. Philion and S. Fidler, "Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d," in *European Conference on Computer Vision*. Springer, 2020, pp. 194–210.
- [27] H. Wang, P. Cai, Y. Sun, L. Wang, and M. Liu, "Learning interpretable end-to-end vision-based motion planning for autonomous driving with optical flow distillation," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 13 731–13 737.
- [28] A. Cui, S. Casas, A. Sadat, R. Liao, and R. Urtasun, "Lookout: Diverse multi-future prediction and planning for self-driving," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 16 107–16 116.
- [29] W. Zeng, S. Wang, R. Liao, Y. Chen, B. Yang, and R. Urtasun, "Ds-dnet: Deep structured self-driving network," in *European conference on computer vision*. Springer, 2020, pp. 156–172.
- [30] A. Sadat, S. Casas, M. Ren, X. Wu, P. Dhawan, and R. Urtasun, "Perceive, predict, and plan: Safe motion planning through interpretable semantic representations," in *European Conference on Computer Vision*. Springer, 2020, pp. 414–430.
- [31] A. Saha, O. Mendez, C. Russell, and R. Bowden, "Translating images into maps," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 9200–9206.
- [32] S. Guadarrama, N. Krishnamoorthy, G. Malkarnenkar, S. Venugopalan, R. Mooney, T. Darrell, and K. Saenko, "Youtube2text: Recognizing and describing arbitrary activities using semantic hierarchies and zero-shot recognition," in *Proceedings of the IEEE international conference on computer vision*, 2013, pp. 2712–2719.
- [33] P. Hanckmann, K. Schutte, and G. J. Burghouts, "Automated textual descriptions for a wide range of video events with 48 human actions," in *European Conference on Computer Vision*. Springer, 2012, pp. 372–380.
- [34] A. Kojima, T. Tamura, and K. Fukunaga, "Natural language description of human activities from video images based on concept hierarchy of actions," *International Journal of Computer Vision*, vol. 50, no. 2, pp. 171–184, 2002.
- [35] M. Rohrbach, W. Qiu, I. Titov, S. Thater, M. Pinkal, and B. Schiele, "Translating video content to natural language descriptions," in *Proceedings of the IEEE international conference on computer vision*, 2013, pp. 433–440.
- [36] S. Chen, J. Chen, Q. Jin, and A. Hauptmann, "Video captioning with guidance of multimodal latent topics," in *Proceedings of the 25th ACM international conference on Multimedia*, 2017, pp. 1838–1846.
- [37] L. Gao, Y. Lei, P. Zeng, J. Song, M. Wang, and H. T. Shen, "Hierarchical representation network with auxiliary tasks for video captioning and video question answering," *IEEE Transactions on Image Processing*, 2022.
- [38] W. Ji and R. Wang, "A multi-instance multi-label dual learning approach for video captioning," *ACM Transactions on Multimedia Computing Communications and Applications*, vol. 17, no. 2s, pp. 1–18, 2021.
- [39] Y. Pan, T. Mei, T. Yao, H. Li, and Y. Rui, "Jointly modeling embedding and translation to bridge video and language," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 4594–4602.
- [40] Y. Pan, T. Yao, H. Li, and T. Mei, "Video captioning with transferred semantic attributes," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 6504–6512.
- [41] R. Pasunuru and M. Bansal, "Multi-task video captioning with video and entailment generation," *arXiv preprint arXiv:1704.07489*, 2017.

- [42] S. Venugopalan, M. Rohrbach, J. Donahue, R. Mooney, T. Darrell, and K. Saenko, "Sequence to sequence-video to text," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 4534–4542.
- [43] S. Venugopalan, H. Xu, J. Donahue, M. Rohrbach, R. Mooney, and K. Saenko, "Translating videos to natural language using deep recurrent neural networks," *arXiv preprint arXiv:1412.4729*, 2014.
- [44] L. Yao, A. Torabi, K. Cho, N. Ballas, C. Pal, H. Larochelle, and A. Courville, "Describing videos by exploiting temporal structure," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 4507–4515.
- [45] H. Yu, J. Wang, Z. Huang, Y. Yang, and W. Xu, "Video paragraph captioning using hierarchical recurrent neural networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 4584–4593.
- [46] B. Zhao, X. Li, X. Lu *et al.*, "Video captioning with tube features," in *IJCAI*, 2018, pp. 1177–1183.
- [47] B. Pan, H. Cai, D.-A. Huang, K.-H. Lee, A. Gaidon, E. Adeli, and J. C. Niebles, "Spatio-temporal graph for video captioning with knowledge distillation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 10870–10879.
- [48] Z. Zhang, Y. Shi, C. Yuan, B. Li, P. Wang, W. Hu, and Z.-J. Zha, "Object relational graph with teacher-recommended learning for video captioning," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 13 278–13 288.
- [49] S. Liu, Z. Ren, and J. Yuan, "Sibnet: Sibling convolutional encoder for video captioning," in *Proceedings of the 26th ACM international conference on Multimedia*, 2018, pp. 1425–1434.
- [50] T. Jin, S. Huang, M. Chen, Y. Li, and Z. Zhang, "Sbat: Video captioning with sparse boundary-aware transformer," *arXiv preprint arXiv:2007.11888*, 2020.
- [51] W. Wang, H. Bao, L. Dong, J. Bjorck, Z. Peng, Q. Liu, K. Aggarwal, O. K. Mohammed, S. Singhal, S. Som *et al.*, "Image as a foreign language: Beit pretraining for all vision and vision-language tasks," *arXiv preprint arXiv:2208.10442*, 2022.
- [52] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 16 000–16 009.
- [53] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 10 012–10 022.
- [54] T. Zhang, X. Chen, Y. Wang, Y. Wang, and H. Zhao, "Mutr3d: A multi-camera tracking framework via 3d-to-2d queries," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 4537–4546.
- [55] M. Han, Z. Zhang, Z. Jiao, X. Xie, Y. Zhu, S.-C. Zhu, and H. Liu, "Scene reconstruction with functional objects for robot autonomy," *International Journal of Computer Vision*, vol. 130, no. 12, pp. 2940–2961, 2022.
- [56] S. Huang, Z. Wang, P. Li, B. Jia, T. Liu, Y. Zhu, W. Liang, and S.-C. Zhu, "Diffusion-based generation, optimization, and planning in 3d scenes," *arXiv preprint arXiv:2301.06015*, 2023.
- [57] Y. Li, Y. Tu, X. Chen, H. Zhao, and G. Zhou, "Distance-aware occlusion detection with focused attention," *IEEE Transactions on Image Processing*, vol. 31, pp. 5661–5676, 2022.
- [58] B. Jin, B. Tian, H. Zhao, and G. Zhou, "Language-guided semantic style transfer of 3d indoor scenes," in *Proceedings of the 1st Workshop on Photorealistic Image and Environment Synthesis for Multimedia Experiments*, 2022, pp. 11–17.
- [59] P. Li, B. Tian, Y. Shi, X. Chen, H. Zhao, G. Zhou, and Y.-Q. Zhang, "Toist: Task oriented instance segmentation transformer with noun-pronoun distillation," *arXiv preprint arXiv:2210.10775*, 2022.
- [60] X. Chen, H. Zhao, G. Zhou, and Y.-Q. Zhang, "Pq-transformer: Jointly parsing 3d objects and layouts from point clouds," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2519–2526, 2022.
- [61] H. Zhao, M. Lu, A. Yao, Y. Chen, and L. Zhang, "Learning to draw sight lines," *International Journal of Computer Vision*, vol. 128, pp. 1076–1100, 2020.
- [62] H. Zhao, M. Lu, A. Yao, Y. Guo, Y. Chen, and L. Zhang, "Physics inspired optimization on semantic transfer features: An alternative method for room layout estimation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 10–18.
- [63] K. Lin, L. Li, C.-C. Lin, F. Ahmed, Z. Gan, Z. Liu, Y. Lu, and L. Wang, "Swinbert: End-to-end transformers with sparse attention for video captioning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 17949–17958.
- [64] J. Janai, F. Güney, A. Behl, A. Geiger *et al.*, "Computer vision for autonomous vehicles: Problems, datasets and state of the art," *Foundations and Trends® in Computer Graphics and Vision*, vol. 12, no. 1–3, pp. 1–308, 2020.
- [65] A. Tampuu, T. Matiisen, M. Semikin, D. Fishman, and N. Muhammad, "A survey of end-to-end driving: Architectures and training methods," *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- [66] L. Yuan, X. Gao, Z. Zheng, M. Edmonds, Y. N. Wu, F. Rossano, H. Lu, Y. Zhu, and S.-C. Zhu, "In situ bidirectional human-robot value alignment," *Science robotics*, vol. 7, no. 68, p. eabm4183, 2022.
- [67] M. Edmonds, F. Gao, H. Liu, X. Xie, S. Qi, B. Rothrock, Y. Zhu, Y. N. Wu, H. Lu, and S.-C. Zhu, "A tale of two explanations: Enhancing human trust by explaining robot behavior," *Science Robotics*, vol. 4, no. 37, p. eaay4663, 2019.
- [68] J. Kim, A. Rohrbach, T. Darrell, J. Canny, and Z. Akata, "Textual explanations for self-driving vehicles," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 563–578.
- [69] X. Chen, T. Liu, H. Zhao, G. Zhou, and Y.-Q. Zhang, "Cerberus transformer: Joint semantic, affordance and attribute parsing," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 19 649–19 658.
- [70] C. Feichtenhofer, A. Pinz, and A. Zisserman, "Detect to track and track to detect," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 3038–3046.
- [71] D. Frossard, E. Kee, and R. Urtasun, "Deepsignals: Predicting intent of drivers through visual signals," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 9697–9703.
- [72] W. Luo, B. Yang, and R. Urtasun, "Fast and furious: Real time end-to-end 3d detection, tracking and motion forecasting with a single convolutional net," in *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2018, pp. 3569–3577.
- [73] S. Casas, W. Luo, and R. Urtasun, "Intentnet: Learning to predict intention from raw sensor data," in *Conference on Robot Learning*. PMLR, 2018, pp. 947–956.
- [74] Z. Liu, J. Ning, Y. Cao, Y. Wei, Z. Zhang, S. Lin, and H. Hu, "Video swin transformer," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 3202–3211.
- [75] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [76] 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 of the Association for Computational Linguistics*, 2002, pp. 311–318.
- [77] S. Banerjee and A. Lavie, "Meteor: An automatic metric for mt evaluation with improved correlation with human judgments," in *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, 2005, pp. 65–72.
- [78] C.-Y. Lin and F. J. Och, "Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics," in *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL-04)*, 2004, pp. 605–612.
- [79] R. Vedantam, C. Lawrence Zitnick, and D. Parikh, "Cider: Consensus-based image description evaluation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 4566–4575.
- [80] F. Yu, H. Chen, X. Wang, W. Xian, Y. Chen, F. Liu, V. Madhavan, and T. Darrell, "Bdd100k: A diverse driving dataset for heterogeneous multitask learning," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 2636–2645.
- [81] J. Carreira, E. Noland, A. Banki-Horvath, C. Hillier, and A. Zisserman, "A short note about kinetics-600," *arXiv preprint arXiv:1808.01340*, 2018.
- [82] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," in *Conference on robot learning*. PMLR, 2017, pp. 1–16.