This WACV 2023 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.

# **DRAMA: Joint Risk Localization and Captioning in Driving**

Srikanth Malla<sup>1</sup> \*

Chiho Choi<sup>1\*\*</sup>

Isht Dwivedi<sup>†</sup>

Joon Hee Choi‡ Jiachen Li ‡‡

### Abstract

Considering the functionality of situational awareness in safety-critical automation systems, the perception of risk in driving scenes and its explainability is of particular importance for autonomous and cooperative driving. Toward this goal, this paper proposes a new research direction of joint risk localization in driving scenes and its risk explanation as a natural language description. Due to the lack of standard benchmarks, we collected a large-scale dataset, DRAMA (Driving Risk Assessment Mechanism with A captioning module), which consists of 17,785 interactive driving scenarios collected in Tokyo, Japan. Our DRAMA dataset accommodates video- and object-level questions on driving risks with associated important objects to achieve the goal of visual captioning as a free-form language description utilizing closed and open-ended responses for multi-level questions, which can be used to evaluate a range of visual captioning capabilities in driving scenarios. We make this data available to the community for further research. Using DRAMA, we explore multiple facets of joint risk localization and captioning in interactive driving scenarios. In particular, we benchmark various multi-task prediction architectures and provide a detailed analysis of joint risk localization and risk captioning. The data set is available at https://usa.honda-ri.com/drama

### 1. Introduction

Situational awareness is an important requirement to achieve high level automation in intelligent vehicles. An important aspect of situational awareness for intelligent mobility is the ability to create an explainable network for the perception of risk from the view-point of the driver, and to establish methods to communicate those risks through voice commands for cooperative and effective driving [16]. In this sense, conveying what the network understood as perceived risk through caption generation can serve two functions: an explainable model for identifying risk agents in the scene,

tjh4.choi@samsung.com, Samsung Semiconductor
ttjiachen\_li@stanford.edu, Stanford University

as well as the ability to generate linguistic descriptions of the scene to communicate those risks through an AI agent.

Important object identification is becoming a key ingredient in autonomous driving and driving assistant systems for safety-critical tasks such as behavior prediction, decision making, and motion planning. One set of approaches have been training the model in a self-supervised manner in addition to their original tasks (e.g., trajectory prediction or steering control), while other approaches have been directly training their model using the human annotated object-level importance. Although explainability of the network predictions is of particular importance, none of existing methods characterizes captioning about identified important objects in the form of a natural language description. Addressing this need, our problem formulation aims to jointly supervise the explainability that is consistent with identification of important objects.

In the absence of an appropriate dataset for this task, we introduce a new dataset to the community to facilitate the study of producing free-form captions explaining the reasoning behind the designation of important objects. A natural language description interprets the interaction of an important object with the ego-vehicle and reasons about its risk observed from ego-driver's perspective in driving scenarios. Our setting of interest and our new dataset, named Driving Risk Assessment Mechanism with A captioning module (DRAMA), are captured through question and answering, seeking closed and open-ended video- and objectlevel attributes from sequence and motion observations.

This dataset provides a basis for a comprehensive exploration of (i) identifying important objects that may influence the future behavior of ego-vehicle; and (ii) captioning about why the system holds its beliefs about future risks of these objects. We designed DRAMA with an explicit goal of visual captioning, particularly in driving scenarios. Each video clip depicts a ego-driver's behavioral response to perceived risk that activates braking of the vehicle. We provide a detailed analysis of such events through questions that are made by elementary operations (what, which, where, why, how) and through answers that represent videoand object-level attributes with a closed and open-ended response, which are utilized to create a free-form caption. Our design choices would encourage in-depth understanding of risks in driving scenarios and allow us to focus on the cap-

<sup>\*</sup>srikanth@kinetic.auto, Kinetic Automation

<sup>\*\*</sup>chihochoi@outlook.com, Samsung Semiconductor

tidwivedi@honda-ri.com, Honda Research Institute

<sup>1</sup> work was majorly done at Honda Research Institute, USA

tioning of these events. To this end, we present algorithmic baselines to characterize the captioning of perceived risks associated with important agents in various multi-tasking settings. The main contributions of the paper are summarized as follows:

- We introduce a novel dataset that provides linguistic descriptions (with the focus on reasons) of driving risks associated with important objects and that can be used to evaluate a range of visual captioning capabilities in driving scenarios.
- We present a data annotation strategy that accommodates the unique characteristics of video- and objectlevel attributes obtained by closed and open-ended responses.
- We provide a new research direction of joint risk localization in driving scenes and its caption as a natural language description and benchmark several baseline models on our DRAMA dataset for these tasks.

# 2. Related Work

Important Object Identification Important agent identification methods can be broadly classified as explicit learning and implicit learning. Explicit learning methods like [40, 11] formulate the agent importance estimation as a binary classification problem. These types of models are typically trained by standard supervised learning. The approaches proposed in [1, 30, 34] learn to imitate human gaze behavior and predict a pixel-level attention map which is used as a proxy for risk. Implicit learning methods proposed in [24, 17, 33, 22] are trained to perform a related task such as trajectory prediction or steering control. Intermediate network activation values are interpreted as perceived risk in these cases. More specifically, the model in [24] is trained to perform future trajectory prediction using self-attention on all the agents in the scene. Self-attention activation outputs are used as agent importance estimates. Causal intervention approach proposed in [22] assumes the counter-examples have fixed ground truth. The approaches in [17, 33] are trained to predict steering control and pixellevel attention heat maps which are used to estimate agent importance. However, none of these methods addresses the reasoning about the model decision nor provides description in the form of natural language, which hinders humans' understanding and the interpretability of safety-critical autonomous driving (AD) and advanced driver-assistance systems (ADAS).

**Vision and Language Grounding** There exist two major categories in vision and language grounding: object retrieval from captions and dense image captioning. In the former, the algorithms such as [9, 13] detect a bounding box in the image given the caption description. Such

models for grounding of the object using natural language have been trained with either a fully supervised [29, 7] or weakly supervised [41] manner using Flickr30K entities dataset [26]. In the latter, dense image captioning methods [15, 39, 37] find several region proposals and predict the caption associated with each proposal using the Visual Genome dataset [20]. Some other works apply visual attention to the image for the task of caption generation [35] or end-to-end learning of driving actions [19]. Our problem setting is closely related to dense image captioning as the language description is generated while detecting the corresponding object. However, our scope is specifically focused on localizing risks in driving scenes where captions provide the reason behind identified important objects.

**Datasets** There exist multiple datasets for accident localization [6, 40, 10] or anomaly detection [38] or causal object localization [28] in driving scenes. Although these benchmarks can be used to understand and predict accidents and related risk attributes, they are inherently incapable of explaining the reason behind important object identification. In-depth analyses of visual explanations can be acquired with the help of linguistic descriptions that explain the perceived risk through a free-form caption. In this view, our DRAMA dataset has unique functions of explanations in the form of a natural language description that is associated with risk localization.

Few driving datasets are aimed to provide commands written in natural language for autonomous vehicles. Talk2Car [9] adds free-form captions on top of the nuScenes dataset [5] to guide the future path of the autonomous car using referred objects in the scene. BDD-X [18] and BDD-OIA [36] has free-form and closed-form explanations respectively, as reasons for ego-vehicle actions. Similarly, HAD [19] provides the vehicle with natural language commands to generate salient maps learned from driver's gaze information. However, these driving language datasets are infeasible for evaluating visual explanations of perceived risks: (i) a localized object is neither a risk nor an important object from the ego-perspective, but simply the object that the caption is referring to; and (ii) labels of these datasets are insufficient to address visual explanations in driving scenes. In [9], captions are generated from a still image without considering temporal relationships against ego-vehicle. Also, per-pixel saliency in [19] is not easily mapped to important objects as driver gaze is often unrelated to the driving. In contrast, DRAMA accommodates ego-driver's behavioral response to external events and explores video- and object-level semantic understanding obtained from video observations.

Our annotation schema of question and answering is similar to the concept of VQA [3, 20]. However, DRAMA specifically addresses visual explanations about driving risks associated with important objects, which cannot be



Figure 1: Example Scenes from DRAMA Dataset, more examples including VQA are available in the supplementary.

discovered using existing dataset: (i) DRAMA facilitates in-depth analysis of driving scenes with annotations that contain domain-specific vocabularies like lane change, right lane, cut-in, traffic congestion, parked vehicle, etc.; and (ii) DRAMA is designed to show why the scenario is risky and what interactions are made as a free-form language description.

# 3. DRAMA Dataset

DRAMA is a dataset that provides visual reasoning of driving risks associated with important objects and that can be used to evaluate a range of visual reasoning capabilities from driving scenarios. We achieve this goal by generating video- and object-level questions and answers based on elementary operations (what, which, where, why, how) that represent object types, visual and motion attributes, object locations, or relationships of objects against the egovehicle. Figure 2 provides an overview of the main components of DRAMA, which we considered for generating the annotations. In the following, we introduce our DRAMA dataset and describe its unique characteristics in detail. We first outline the data collection procedure including sensor configurations, raw data details, annotation schema, quality, and privacy in Section 3.1. Next, we provide a statistical analysis of the dataset in Section 3.2 and highlight the uniqueness by comparing with publicly available datasets in Section 3.3. Example scenarios of the DRAMA dataset are shown in Figure 1.

#### 3.1. Dataset Creation

We recorded 91 hours of videos from urban roads in Tokyo, Japan using SEKONIX SF332X-10X video camera (30HZ frame rate, 1928 × 1280 resolution and 60° H-FOV) and GoPRO Hero 7 camera (60HZ frame rate, 2704 × 1520 resolution and 118.2° H-FOV). The videos are synchronized with the Controller Area Network (CAN) signals and Inertial Measurement Unit (IMU) information. We filtered out these videos based on the ego-driver's behavioral response to external situations or events, which activate braking of the vehicle. As a result, 17,785 interactive driving scenarios are created where each scenario is 2 seconds in duration. The annotators are requested to watch each scenario as observations and instructed to generate labels as described in the following paragraph.

The annotations of DRAMA are designed for targeting identification of important objects and generation of associated attributes and captions based on direct / indirect influences on the ego-car's motion. To obtain data in this form, we developed an annotation schema as follows:

**Video-level Q/A** Questions are formatted to address risk and scene attributes as well as ego-behavioral functions such as ego-intention and behavior suggestions. Figure 2 (left) illustrates these attributes in DRAMA. Some questions ask whether risks are perceived or what the driver's potential action is in order to ease the driving risk. The valid answers are closed-form such as a boolean format (*i.e.*, yes or no) or single choice from multiple options (*e.g.*, stop, slow down, start moving, ...).



Figure 2: Annotation Schema and example labels for the DRAMA Dataset

**Object-level Q/A** If the scenario is determined risky, the annotators proceed to a next step, object-level Q/A. It includes elementary operations (what, which, where, why, how) that can structure a question allowing a free-form or open-ended response. In addition to single choice answers, more complex reasoning is allowed to represent the high-level semantic understanding from observations. Figure 2 (right) shows example answers that correspond to our elementary operations of visual reasoning such as querying object types (what), visual and motion attributes (which), location and motion direction (where), reasoning about risks or description of interactions (why), and effect of relation-ships (how).

For more details on privacy and quality consistency annotation procedure, please check supplementary material.

# **3.2. Dataset Statistics**

Figure 3a shows the distribution of labels obtained from video-level Q/A. Given the filtering of raw videos based on the activation of vehicle braking using CAN, majority of scenarios (17,066, 95.95%) are labeled risky with the perception of important objects. The reasoning about risks was made during interactions of the important object when the ego-vehicle goes straight (94.5%), turns left (2.6%), or turns right (2.9%). Each scenario was further categorized into three groups based on interactive environments: wider roads (51.5%), intersections (32.6%), and narrow streets (15.9%).

Figure 3b presents the distribution of object-level Q/A attributes. The object that is recognized important was a vehicle (71.9%), pedestrian or cyclist (19.6%), and infrastructure (8.5%). The visual attribute of these important objects can be generally answered with open-ended responses such as a red sedan vehicle, a tall person wearing blue shirt, red traffic light, etc., which results in 35,038 visual attributes

from 17,066 objects. Similarly, motion attributes, location, motion directions, and reasoning about agent behaviors also have free-form values. In particular, the descriptions of reasoning include 992 unique words with total occurrences of 306,708 times. The number of unique words for vehicle, pedestrian or cyclist, and infrastructure is 533, 608 and 268, respectively. Also, 112 unique words are commonly shown across all object classes, which consist of colors (white, black, ...), locations (front, right, left), articles (the, an), prepositions (from, in, of, at), auxiliary verbs (is, are), adverbs (ahead, ...), nouns (street, road, car, ego, ...), conjunctions (because, and ), and actions (stopped, parked). The distribution of the words is shown in Figure 4. More details are provided in the supplementary.

# 3.3. Dataset Analysis

Table 1 shows the comparison of the DRAMA dataset to existing benchmarks. HDD [27], Epic Fail (EF) [40], Street Accident (SA) [6], and DoTA [38] are object localization datasets that benchmark risk localization or anomaly detection. HAD [19], BDD-X [18], BDD-OIA [36] are captioning datasets in driving that address advises and reasons for ego-vehicle actions in natural language descriptions. Although these datasets are similar to DRAMA to some extent, we jointly address risk localization with visual reasoning of driving risks with a free-form language description. Talk2Car (T2C) [9] is the image dataset that provides a caption associated with a bounding box. However, their object labels are not related to driving risks, and more importantly the caption is generated in a way that requests for an action to a car. In contrast, DRAMA provides important object labels with captions describing their risk from the ego-car perspective, while considering spatio-temporal relationships from videos. We directly compare our example scenario with that of [9] in Figure 5. Our caption labels are





(a) 40 common words across all labels

(b) Word distribution per class



1.4x more captions with 1.6x longer in length with respect to T2C. Unlike these benchmark datasets, DRAMA also delivers the functionality of visual question and answering to evaluate a range of visual reasoning capabilities from driving scenarios.

#### 4. Methodology

We introduce a model to address risk localization and its reasoning, which consist of an encoder and decoder as in Figure 6. The details of model architecture are shown in the supplementary material.

**Encoder** Given an image  $I_t$  and an optical flow image  $F_t$  at time step t, the encoder extracts visual and optical flow features. The optical flow image  $F_t$  is generated by Flownet 2.0 [14] with two image frames  $I_t$  and  $I_{t-1}$  to capture the motion of objects and scene. For notational brevity, we drop t in the following sections. To encode visual feature and optical flow feature, we use a ResNet-101 model [12] pre-trained on Imagenet [21] as a backbone network. Both features are extracted from two layers prior to the original ResNet-101 structure, and we add the adaptive average pooling layer to get a fixed embedding size of 14x14x2048 as outputs. We further concatenate these two features to generate the output feature E.

**Decoder** The decoder is trained to find a bounding box  $(\hat{y}_b)$  as an important object and to generate a corresponding caption description  $(\hat{y}_c)$ . In our proposed model, we first

predict a caption  $\hat{y}_c$  using an LSTM-based recurrent model. The hidden states of LSTM,  $H = [h_1, \ldots, h_M]$ , are used as the embedding for the self-attention [31]. The resulting new feature  $H_b = [h_{b1}, \ldots, h_{bM}]$  is obtained by  $h_{bi} =$  $self attention(HW_i^Q, HW_i^K, HW_i^V)$ , where  $W^Q, W^K$ and  $W^V$  are parameter matrices for queries, keys and values, respectively. Lastly, we find a bounding box  $\hat{y}_b$  using an important object localization model. In the experiments, the caption prediction model follows the decoder architecture of SAT [35], which is a LSTM model with the softattention mechanism. Also, we used a basic multi-layer perceptron (MLP) for the important object localization model.

**Loss Function** The loss function for joint risk localization and reasoning is defined as follows:

$$L_b = \sum_{j=1}^{N} \left\| \hat{y}_b^{(j)} - y_b^{(j)} \right\|_2;$$
(1)

$$L_{c} = \sum_{j=1}^{N} \left[ \sum_{k=1}^{M} ce\left(\hat{y}_{ck}^{(j)}, y_{ck}^{(j)}\right) + \lambda \sum_{i=1}^{P} \left( 1 - \sum_{k=1}^{M} \alpha_{ki}^{(j)} \right)^{2} \right],$$
(2)

where  $^{(j)}$  denotes the *j*-th data and the right term of  $L_c$  is a doubly stochastic attention regularizer adapted from SAT [35].  $\alpha_{ki}$  denotes an element in the attention map in the caption prediction, M and P are the caption length and the number of pixels in the attention map, respectively, and

	Risk Localization only			Localization and Captioning		Captioning only			
Dataset	HDD [28]	EF [40]	SA [6]	DoTA [38]	DRAMA (ours)	T2C [9]	HAD [19]	BDD-X [18]	BDD-OIA [36]
Risk localization	1	1	1	1	1	no	no	no	no
Risk captions	no	no	no	no	1	no	1	1	1
# of scenarios	137	3,000	1,733	4677	17,785	9,217	5,744	6,984	11,303
# of captions	-	-	-	-	17,066	11,959	22,366	26,534	27,265
Avg caption length	-	-	-	-	17.97	11.01	11.05	$8.90^{1}$	6.81 <sup>1</sup>
VQA	no	no	no	no	1	no	no	no	no
# of questions	-	-	-	-	77,639	-	-	-	-
# of answers	-	-	-	-	102,830	-	-	-	-
Environment attribute	1	no	no	no	1	no	no	no	no
Reason / Why attribute	1	no	no	no	1	no	no	1	1
Scene-level risk attribute	1	1	1	1	1	no	no	no	no
Freeform captions	no	no	no	no	1	no	1	1	no

Table 1: Comparison of DRAMA with public benchmark driving datasets.



I see jaime's truck. It is the one in white parked on the left. Pull over here so I can have a chat with him. There is a white van stopped on the left side of the road because it is waiting to turn onto the ego lane.

(c) DRAMA VQA Video-level VQA: Risk: Yes Ego-intention: Straight Scene-Classifier: Urban Road Object-level VQA What (object): Vehicle Whato (visual): White, Van Which (motion): Stopped Where (location): left side of the road Where (motion): into ego-lane Why (reason): waiting to turn onto

the ego-lane

Figure 5: Comparison of an example scenario between (a) Talk2Car [9] and (b) our DRAMA dataset. Both scenarios annotate a white vehicle on the left as an object label (green bounding box) and provide an associated caption. (c) shows a snippet of video- and object-level attributes of DRAMA, corresponding to the scenario of (b).

 $\lambda$  is a hyperparameter. Then, we use a multi-task loss [8] for combining  $L_b$  and  $L_c$ , which results in

$$L_{total} = \frac{1}{\sigma_1^2} L_b + \frac{1}{\sigma_2^2} L_c + \log(\sigma_1 \sigma_2),$$
(3)

where  $\sigma_1$  and  $\sigma_2$  denote the relative weights of the losses  $L_b$ and  $L_c$ , respectively. The multi-task loss function considers the homoscedastic uncertainty of each task by training the weights  $\sigma_1$  and  $\sigma_2$ .

### 5. Experiments

#### 5.1. Model Comparison

We evaluate the performance of our model on the new DRAMA dataset, by comparing with various baseline models. For the fair comparison, all baselines use the same backbone networks; pre-trained ResNet-101 [12] for the encoder, the SAT [35] structure for caption prediction (CP), and MLP for important object localization (IOL). Since our model is conditioned on captioning for localization, we named LCP (localization with captioning prior). Our baseline models are described as follows:

**ResNet-101** [12] This is a strong baseline model that identifies an important object only, which is trained with Equation (1). As mentioned earlier, our structure for the localization task is built from the original ResNet model by replacing its last two layers by adaptive average pooling and MLPs as the decoder. Therefore, this model provides an insight on the efficacy of an additional optical flow input as well as our total loss function in Equation (3).

**SAT** [35] This baseline model only generates a caption description. The model takes both an image and optical flow as input and have the loss function of Equation (2). We compare this baseline with LCP to demonstrate our captioning capability under joint training of risk localization and reasoning.

**Independent captioning and localization (ICL)** The ICL model simultaneously finds the important object  $\hat{y}_b$  and caption description  $\hat{y}_c$  independently from one another. This model is trained using the dual task loss function in Equation (3), same as our LCP model.

**Captioning with localization prior (CLP)** In opposition to LCP, this baseline model has the decoder architecture that generates captioning using a localization prior. We add a region of interest (ROI) pooling to improve the model's captioning capability. The ROI pooling produces the im-

<sup>&</sup>lt;sup>1</sup>reported average caption length is without ego action for fair comparison with our dataset, otherwise it is 13.92 for BDD-X and 8.36 for BDD-OIA.



Figure 6: DRAMA Architecture: the encoder and decoder modules

portant object feature  $E_a$  from the encoder output E. Then, we concatenate  $E_a$  with the output of IOL, and the concatenated features are used as the input of the caption generation. This model is similar to [15] with a small variation of skip connection of the encoder output E, to get global information. This model is also trained using the dual task loss function in Equation (3).

We evaluate the effectiveness of optical flow information by creating additional baselines using LCP, ICL and CLP models without optical flow stream. Furthermore, we generate additional LCP without the multi-task loss function, which validates our effort of balancing the multi-task training.

#### 5.2. Experimental Setting and Metrics

We split the DRAMA dataset into 70% train, 15% validation and 15% test. The model parameters are chosen on the validation set at training time, and we report the performance of this model tested on the test set. The captioning performance is evaluated using the standard metrics such as BLEU-1 (B1) [25], BLEU-4 (B4) [25], METEOR (M) [4], ROGUE (R) [23], CIDER (C) [32], and SPICE (S) [2]. We use Mean-IOU (Intersection Over Union) and accuracy for IOU>0.5 for the evaluation of important agent localization. IOU is a standard metric in object detection, dividing the area of overlap by the area of union between the predicted bounding box and the ground-truth bounding box. We provide other experimental settings in the supplementary.

#### 5.3. Results

**Quantitative comparison** Table 2 shows the quantitative results of our model and baseline models on the DRAMA dataset. First, we can see that adding the optical flow information is a positive effect in captioning across all the decoder baselines including LCP, ICL and CLP. The reason is the words related to motion such as moving, stopped, etc. can be captured by the optical flow information, but not by a single image. Since the video dataset can only use optical flow images, our DRAMA dataset has more benefits than the image dataset such as the Talk2Car [9] dataset. Next, our LCP model performs better than ICL and CLP.

baselines in important object identification. It makes sense that LCP performs well in the localization due to an additional self-attention module. Note that the captioning performance of LCP is higher or similar to that of CLP and better than those of ICL and SAT. It implies that better localization derived by self-attention influenced to finding a better multi-task weighting in our loss function. Although the ROI pooling in the CLP model improves the captioning comparing to ICL and SAT, it seems to negatively affect to risk localization. Additionally, it shows dual task loss helps to improve the overall performance of LCP (LCP vs LCP without DTL).

We compare our model to the single task baselines, SAT and ResNet-101, in detail with respect to the type of objects. We use the CIDER and Mean-IOU metrics for the comparison. Table 3 shows that our model performs better than SAT for all types of objects in the captioning. Also, important object identification of our method is better than that of the ResNet-101 model for most of object types. It demonstrates that both tasks are highly correlated and can benefit being tackled together.

**Qualitative comparison** Figure 7 shows the ground truth and our prediction in the cases where the important objects are pedestrians (a,g), vehicles (b,c,e,f) and infrastructure (d,h). Figure 7a and 7b show the model learns a good reasoning such as "because of another pedestrian" and "because of green traffic light" even though the ground truth does not have. In Figure 7c and 7e, the reasoning "because of traffic congestion ahead" is correctly predicted by our model. In Figure 7f, the prediction correctly describes the object in a different way than the ground truth, e.g., "turning right in the middle of the intersection" vs. "cutting in from the right side of the intersection." Figure 7g and 7h show our model fails to find the ground-truth important object and caption description. Although our model identifies the object closer to the ego-vehicle rather than the object heading towards the ego-vehicle in Figure 7g, it is also considered as a driving risk observed in the scene. We believe such an issue can be addressed by generating multi-modal outputs, which is our future research direction. In Figure 7h, the red traffic light is closer and more important than the vehi-



Figure 7: Qualitative results of the DRAMA model: our prediction in red and ground truth in green. More detailed success and failure cases are discussed in the supplementary material.

	Captioning					Important Object Localization		
Method	B1 ↑	<b>B</b> 4 ↑	M↑	R ↑	C ↑	S ↑	Mean-IOU ↑	Acc (IOU>0.5) $\uparrow$
ResNet-101 [12]	-	-	-	-	-	-	0.600	0.670
SAT [35]	0.740	0.531	0.386	0.694	3.583	0.544	-	-
ICL without OF	0.731	0.518	0.377	0.679	3.468	0.522	0.553	0.617
ICL	0.739	0.539	0.385	0.693	3.678	0.534	0.533	0.593
CLP without OF [15]	0.744	0.544	0.386	0.700	3.582	0.524	0.518	0.587
CLP	0.744	0.546	0.388	0.701	3.670	0.540	0.520	0.581
LCP without OF	0.721	0.520	0.379	0.688	3.501	0.536	0.597	0.667
LCP without DTL	0.740	0.543	0.389	0.703	3.668	0.548	0.569	0.640
LCP (Ours)	0.739	0.547	0.391	0.700	3.724	0.560	0.614	0.684

Table 2: Quantitative evaluation comparing our model LCP (localization with captioning prior) with ICL (independent captioning and localization) and CLP (captioning with localization prior). OF and DTL denotes optical flow and dual task loss function, respectively.

		C↑	Mean-IOU ↑		
Object	SAT	LCP (ours)	ResNet-101	LCP (ours)	
Vehicle	3.588	3.717	0.714	0.737	
Pedestrian	2.100	2.360	0.266	0.286	
Cyclist	1.630	1.912	0.276	0.271	
Infrastructure	6.216	6.480	0.091	0.109	

Table 3: LCP is compared with SAT [35] and ResNet-101 [12] with respect to object types.

cle, but our model identifies the vehicle mainly because the model is incapable of understanding the physical distance between objects. Extension to incorporate such geometric information is our next plan.

#### 6. Conclusion

We introduced a new research direction of joint risk localization in driving scenes and its explanations as a natural language description. To address this problem, we collected a novel driving dataset, DRAMA, which is annotated with video- and object level questions on important objects as well as an interpretation of the scene with respect to the interaction observed from the ego-driver's perspective. As a result, DRAMA facilitates the evaluation of a range of visual explanation capabilities from driving scenarios. We designed various models for the purpose of benchmarking jointly supervised frameworks that localize a risk and its linguistic explanations as a natural language description. Extensive ablative study on different architectures demonstrated that important object localization using captioning as a prior showed the improved performance against other baseline models.

### References

[1] Stefano Alletto, Andrea Palazzi, Francesco Solera, Simone Calderara, and Rita Cucchiara. Dr (eye) ve: a dataset for attention-based tasks with applications to autonomous and assisted driving. In *Proceedings of the ieee conference on computer vision and pattern recognition workshops*, pages 54–60, 2016.

- [2] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. Spice: Semantic propositional image caption evaluation. In *European conference on computer vision*, pages 382–398. Springer, 2016.
- [3] 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*, pages 2425– 2433, 2015.
- [4] 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 intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72, 2005.
- [5] 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*, pages 11621–11631, 2020.
- [6] Fu-Hsiang Chan, Yu-Ting Chen, Yu Xiang, and Min Sun. Anticipating accidents in dashcam videos. In Asian Conference on Computer Vision, pages 136–153. Springer, 2016.
- [7] Kan Chen, Rama Kovvuri, and Ram Nevatia. Query-guided regression network with context policy for phrase grounding. In *Proceedings of the IEEE International Conference* on Computer Vision, pages 824–832, 2017.
- [8] Roberto Cipolla, Yarin Gal, and Alex Kendall. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition, pages 7482– 7491, 2018.
- [9] 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.
- [10] Jianwu Fang, Dingxin Yan, Jiahuan Qiao, and Jianru Xue. Dada: A large-scale benchmark and model for driver attention prediction in accidental scenarios. *arXiv preprint arXiv:1912.12148*, 2019.
- [11] Mingfei Gao, Ashish Tawari, and Sujitha Martin. Goaloriented object importance estimation in on-road driving videos. In 2019 International Conference on Robotics and Automation (ICRA), pages 5509–5515. IEEE, 2019.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceed-ings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [13] Ronghang Hu, Huazhe Xu, Marcus Rohrbach, Jiashi Feng, Kate Saenko, and Trevor Darrell. Natural language object retrieval. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 4555–4564, 2016.

- [14] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2462–2470, 2017.
- [15] Justin Johnson, Andrej Karpathy, and Li Fei-Fei. Densecap: Fully convolutional localization networks for dense captioning. In *Proceedings of the IEEE conference on computer* vision and pattern recognition, pages 4565–4574, 2016.
- [16] Helen Keyes, Antony Whitmore, Stanislava Naneva, and Daragh McDermott. The priming function of in-car audio instruction. *Quarterly Journal of Experimental Psychology*, 72(3):643–650, 2019.
- [17] Jinkyu Kim and John Canny. Interpretable learning for selfdriving cars by visualizing causal attention. In *Proceedings* of the IEEE international conference on computer vision, pages 2942–2950, 2017.
- [18] Jinkyu Kim and et al. Textual explanations for self-driving vehicles. In *ECCV*, pages 563–578, 2018.
- [19] Jinkyu Kim, Teruhisa Misu, Yi-Ting Chen, Ashish Tawari, and John Canny. Grounding human-to-vehicle advice for self-driving vehicles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10591–10599, 2019.
- [20] 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 vision*, 123(1):32–73, 2017.
- [21] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. Advances in neural information processing systems, 25:1097–1105, 2012.
- [22] Chengxi Li and et al. Who make drivers stop? towards driver-centric risk assessment: Risk object identification via causal inference. In 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 10711–10718. IEEE, 2020.
- [23] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004.
- [24] Srikanth Malla, Behzad Dariush, and Chiho Choi. Titan: Future forecast using action priors. In *Proceedings of* the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 11186–11196, 2020.
- [25] 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.
- [26] Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proceedings of the IEEE international conference on computer* vision, pages 2641–2649, 2015.

- [27] Vasili Ramanishka, Abir Das, Jianming Zhang, and Kate Saenko. Top-down visual saliency guided by captions. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 7206–7215, 2017.
- [28] Vasili Ramanishka and et al. Toward driving scene understanding: A dataset for learning driver behavior and causal reasoning. In *CVPR*, pages 7699–7707, 2018.
- [29] Anna Rohrbach, Marcus Rohrbach, Ronghang Hu, Trevor Darrell, and Bernt Schiele. Grounding of textual phrases in images by reconstruction. In *European Conference on Computer Vision*, pages 817–834. Springer, 2016.
- [30] Ashish Tawari, Praneeta Mallela, and Sujitha Martin. Learning to attend to salient targets in driving videos using fully convolutional rnn. In 2018 21st International Conference on Intelligent Transportation Systems (ITSC), pages 3225– 3232. IEEE, 2018.
- [31] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. 2017.
- [32] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 4566–4575, 2015.
- [33] Dequan Wang, Coline Devin, Qi-Zhi Cai, Fisher Yu, and Trevor Darrell. Deep object-centric policies for autonomous driving. In 2019 International Conference on Robotics and Automation (ICRA), pages 8853–8859. IEEE, 2019.
- [34] Ye Xia, Danqing Zhang, Jinkyu Kim, Ken Nakayama, Karl Zipser, and David Whitney. Predicting driver attention in critical situations. In *Asian conference on computer vision*, pages 658–674. Springer, 2018.
- [35] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057. PMLR, 2015.
- [36] Yiran Xu and et al. Explainable object-induced action decision for autonomous vehicles. In *CVPR*, pages 9523–9532, 2020.
- [37] Linjie Yang, Kevin Tang, Jianchao Yang, and Li-Jia Li. Dense captioning with joint inference and visual context. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 2193–2202, 2017.
- [38] Yu Yao, Xizi Wang, Mingze Xu, Zelin Pu, Ella Atkins, and David Crandall. When, where, and what? a new dataset for anomaly detection in driving videos. *arXiv preprint arXiv:2004.03044*, 2020.
- [39] Guojun Yin, Lu Sheng, Bin Liu, Nenghai Yu, Xiaogang Wang, and Jing Shao. Context and attribute grounded dense captioning. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 6241– 6250, 2019.
- [40] Kuo-Hao Zeng, Shih-Han Chou, Fu-Hsiang Chan, Juan Carlos Niebles, and Min Sun. Agent-centric risk assessment: Accident anticipation and risky region localization. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 2222–2230, 2017.

[41] Fang Zhao, Jianshu Li, Jian Zhao, and Jiashi Feng. Weakly supervised phrase localization with multi-scale anchored transformer network. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition, pages 5696– 5705, 2018.