

# Rank2Tell: A Multimodal Driving Dataset for Joint Importance Ranking and Reasoning

Enna Sachdeva<sup>\*1</sup>, Nakul Agarwal<sup>\*1</sup>, Suhas Chundi<sup>2</sup>, Sean Roelofs<sup>2</sup>, Jiachen Li<sup>2</sup>, Mykel Kochenderfer<sup>2</sup>, Chiho Choi<sup>1,†</sup>, Behzad Dariush<sup>1</sup>

<sup>\*</sup>equal contribution

<sup>1</sup>Honda Research Institute USA      <sup>2</sup>Stanford University

## Abstract

*The widespread adoption of commercial autonomous vehicles (AVs) and advanced driver assistance systems (ADAS) may largely depend on their acceptance by society, for which their perceived trustworthiness and interpretability to riders are crucial. In general, this task is challenging because modern autonomous systems software relies heavily on black-box artificial intelligence models. Towards this goal, this paper introduces a novel dataset, Rank2Tell<sup>1</sup>, a multi-modal ego-centric dataset for **Ranking** the importance level and **Telling** the reason for the importance. Using various close and open-ended visual question answering, the dataset provides dense annotations of various semantic, spatial, temporal, and relational attributes of various important objects in complex traffic scenarios. The dense annotations and unique attributes of the dataset make it a valuable resource for researchers working on visual scene understanding and related fields. Furthermore, we introduce a joint model for joint importance level ranking and natural language captions generation to benchmark our dataset and demonstrate performance with quantitative evaluations.*

## 1. Introduction

Effective and accurate understanding of visual scenes is an important prerequisite for safe navigation of autonomous vehicles and advanced driver assistance systems, especially in complex and highly crowded urban scenarios. Despite significant advancement in the development of self-driving and driver-support technologies, public acceptance of these systems remains limited. A survey conducted by Partners for Automated Vehicle Education (PAVE) in 2020 [12] reported that 60% of respondents would be more likely to trust autonomous vehicles if they have a better understanding of the underlying rationale of the models. To im-

prove the transparency of these systems, intelligent vehicles must have the ability to identify critical traffic agents whose behaviors can influence their own decision making [22, 23, 27, 47]. Identifying these important agents allows for a more efficient allocation of computation resources toward predicting the actions of a subset of critical objects and identifying potential risks. To establish trust, the autonomous system must provide human-interpretable reasoning about the important agents in the scene, through voice or visual interfaces.

Accurately identifying important agents within the environment and providing human-interpretable reasoning requires a scene understanding model to capture several essential scene features effectively, including 3D mapping, the semantics of the scene, spatial and temporal relations, agents' importance level, actions, intentions, and attention. Additionally, the ability to reason about important agents in a human-interpretable manner is crucial for capturing the essence of the scene. By integrating these comprehensive features, the model's ability to understand and reason about the scene is greatly enhanced. One key to successfully applying these approaches is the availability of traffic datasets with rich human annotations of object importance and reasoning to address the interpretability and trustworthiness of autonomous systems operating in an interacting environment. However, there currently exists no comprehensive real-world driving dataset that provides all these features.

To address this challenge and facilitate future research, we propose a novel ego-centric, multi-modal dataset for visual scene understanding in urban traffic scenarios. The dataset uses 2D image features and 3D LiDAR features to provide dense semantics, temporal and relational annotations of important agents that influence the ego vehicle's decision making. In addition, it provides diverse natural language explanations to enable reasoning about why a particular agent in a scene is of importance. We aim to improve the transparency and interpretability of the visual scene understanding modules of autonomous systems. The proposed

<sup>†</sup>Now at Samsung Semiconductor Inc.

<sup>1</sup><https://usa.honda-ri.com/rank2tell>

dataset has the potential to assist drivers in conveying important decisions, improve their situational awareness, and warn passengers about potential safety hazards present in the surrounding environment. We also introduce a model that uses multi-modal 2D+3D features to jointly predict the importance level and generate captions of important agents.

The contributions of this paper are as follows. **First**, we introduce the first multi-modal dataset, *Rank2Tell*, for importance level ranking and natural language explanation tasks in urban traffic scenarios. The data is annotated using visual question answering (VQA) that combines video and object-level attributes. **Second**, we propose a model that uses multi-modal features for joint importance level classification and natural language captioning, and also establish a benchmark suite on these tasks. **Third**, we introduce the key features of the dataset that can potentially be used to enhance scene understanding for safety-critical applications.

## 2. Related Work

### 2.1. Object Importance Classification

Identifying important objects in images or videos has attracted a great deal of research interest in various domains. There are several visual saliency studies that have created a pixel-level importance mapping for tasks such as visual question answering [1,45], scene understanding [35], video summarising [15], driver attention [24], and AV steering control [17]. However, these methods do not treat objects as distinct instances. Several studies focus on identifying the single most important object in a scene, such as the most important person in a social situation [25]. Several other studies make a binary classification between “important” and “unimportant” for multiple objects in a scene [13,21,22,48]. However, only one of these studies [22] attempts to couple natural language with importance estimation but this work only provides a description of a single important object. Other studies [29,39] classify agents in the scenes at three different importance levels: high, medium, and low. However, none of these studies explains the underlying reasoning for the importance estimation for a specific object.

### 2.2. Dense Captioning

There has been a large amount of research on generating natural language to describe visual images or videos. The idea of self-attention was proposed in the visual captioning task [40], which has been the basis for many further visual captioning research including [3,10]. Some dense visual captioning methods aim to generate a caption for multiple regions in the same image [16,42,44] using the Visual Genome dataset [20]. The ScanRefer dataset was proposed for object localization and language description tasks [5]. Our problem setting is similar to dense image captioning but we focus on generating explanations for the important ob-

jects in traffic scenarios with spatio-temporal observations.

### 2.3. Datasets

In recent years, many traffic scene datasets have been proposed to stimulate progress in the analysis of the important objects in driving scenes. HDD [32] and KITTI [29] are object localization datasets that benchmark importance localization and anomaly detection. HAD [18], BDD-X [19], BDD-OIA [41] are captioning driving datasets that provide reasons for the ego vehicle’s actions in natural language descriptions. DRAMA [28] provides important object labels with captions from the ego car’s perspective while considering spatio-temporal relationships from videos. While DRAMA [28] is the most relevant dataset to our proposed Rank2Tell, the major differences are shown in Table 1.

## 3. Rank2Tell Dataset

The dataset is collected using an instrumented vehicle equipped with three Point Grey Grasshopper video cameras with a resolution of  $1920 \times 1200$  pixels, a Velodyne HDL-64E S2 LiDAR sensor, and high precision GPS. Additionally, Vehicle Controller Area Network (CAN) data is collected for analyzing how drivers manipulate steering, breaking, and throttle. All sensor data are synchronized and timestamped using ROS and customized hardware and software. The dataset contains a diverse set of traffic scenes captured at intersections in different environments, including urban areas. We selected 116 clips of approximately 20 seconds each, focusing on intersections from several hours of data. These video clips capture both the entering and exiting of the ego vehicle from the intersection.

### 3.1. Annotation Methodology

The important agent identification is subjective in nature and can vary based on factors such as age, gender, and driving experience. To account for this diversity of opinions, each video was annotated by five annotators with varying levels of driving experience and age. The detailed statistics of annotators are introduced in the supplementary material.

We first stitch images from the three cameras, i.e. front, left, and right, which provide a wide FOV (horizontal 134 degrees). Each frame of a video is overlaid with the ego vehicle’s speed, and intention (left/right/straight) while leaving the intersection using an arrow. Annotators annotated every 4th frame of the 10 fps video. To avoid bias in important agent identification due to prior knowledge of other agents’ future intentions, the annotators had access to only 40 historical frames (i.e., 4 seconds of the video) while annotating a certain frame.

We divide our annotation scheme into three parts: *Important Agent Identification*, *Important Agent Localization and Ranking*, and *Captioning*. The annotation schema is shown in Figure 2.





Figure 1. **Overview of Rank2Tell.** Rank2Tell is an ego-centric dataset with three camera images and point cloud features for visual scene understanding in complex urban traffic scenarios. For each scenario, five annotators are asked to identify important objects in the scene with three importance levels: high, medium, and low. They also compose natural language descriptions to explain their reasoning behind the importance ranking of each important object in the scene, which leads to diverse annotations of explanations. In this exemplar scenario, the important objects are indicated by bounding boxes. The importance and natural language annotations of the traffic light and the crossing pedestrian are shown as illustrative examples.

Table 1. Comparison of our proposed dataset with other datasets.


Dataset	Important Agents Identification only						Captioning only			Both		
	HDD [32]	EF [46]	SA [4]	DoTA [43]	KITTI [29]	A-SASS [39]	HAD [18]	BDD-X [19]	BDD-OIA [41]	T2C [9]	DRAMA [28]	Ours
# Scenarios	137	3,000	1733	4,677	-	10	5,744	6,984	11,303	850	17,785	116
# Frames annotated/video	$\geq 1$	$\geq 1$	$\geq 1$	$\geq 1$	-	$\geq 1$	$\geq 1$	$\geq 1$	$\geq 1$	1	1	$\geq 1$
Avg. video duration (in sec)	2700	4	5	-	-	-	20	40	5	20	2	20
Importance Localization	✓	✓	✓	✓	✗	✓	✗	✗	✗	✗	✓	✓
Importance Captioning	✗	✗	✗	✗	✗	✗	✓	✓	✓	✗	✓	✓
Importance Ranking	✗	✗	✗	✗	✓	✓	✗	✗	✗	✗	✗	✓
# Important objects/frame	1	1	$\geq 1$	$\geq 1$	$\geq 1$	$\geq 1$	-	-	-	-	1	$\geq 1$
# Captions/frame	-	-	-	-	-	-	1	1	1	-	1	$\geq 1$
# Captions/object	-	-	-	-	-	-	1	1	1	-	1	$\geq 1$
Avg caption length	-	-	-	-	-	-	11.05	8.90	6.81	11.01	17.97	31.95
RGB	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
LiDAR	✗	✗	✗	✗	✓	✗	✗	✗	✗	✓	✗	✓
3D boxes	✗	✗	✗	✗	✓	✗	✗	✗	✗	✓	✗	✓
Field of View (RGB)	C	C	C	C	C	LCR	C	C	C	LCR	C	LCR
Object Tracking	✗	✗	✗	✗	✓	✗	✗	✗	✗	✗	✗	✓
Reasoning	✓	✗	✗	✗	✗	✗	✗	✓	✓	✗	✓	✓
Free Form Captions	✗	✗	✗	✗	✗	✗	✓	✓	✗	✓	✓	✓

**Important Agent Identification.** The first step is to identify whether there exist important agents in the scene that have the potential to influence the ego vehicle. To this end, we overlay the ego vehicle’s intention (i.e., straight, left, right) and speed obtained using CAN bus data, as shown in Figure 1. Then we overlay this information on the stitched images before passing it on to the annotators,


which provides important context information to annotators about driver behavior and intent which is crucial to identify important agents. This is in contrast with other works in that [29] does not provide this information at all and [28] asks the annotators to label the ego vehicle’s intention along with importance, for which they need to watch the entire clip which is not a realistic setting in the real world because



(a) Sequence of Visual Questions asked from annotators during the annotation



(b) Location Level



(c) Motion Direction

Figure 2. The annotation schema of Rank2Tell dataset.

the driver does not have access to future information. Additionally, while [28] filters raw videos based on the activation of vehicle braking using CAN information, we manually filter data, which is more accurate. The annotators then watch the video and are asked to imagine themselves as the driver and determine if there are important agents present in the scene that may affect their driving.

**Important Agent Localization and Ranking.** The annotators are instructed to localize each agent in the scene that is important to the ego vehicle by creating a bounding box. At a high level, we asked the annotators to identify agents to which the ego vehicle should be attentive for safe driving. Once an object is marked as important, annotators draw a 2D bounding box around that agent, and rank its importance level and relevance they would have given the object in real driving. In this work, we use three levels of importance: *Low, Medium, High*.

The purpose is to handle ambiguity which is hard to

avoid in case of binary categorization, i.e. important or non-important [22], and also to reduce confusion and guesswork which is inevitable in case of a continuous ranking score. Inspired by [29, 39], we posit that two levels of importance (i.e., important and nonimportant) could be overly restrictive and inadequate for handling ambiguous cases. Moreover, [39] shows that multiple levels could aid drivers' situational awareness in real time, with the minimal distraction of multiple important objects. Due to the subjectivity of the task, we used five annotators to label each scene to reduce ambiguity and reach some level of consensus. Since different annotators may perceive importance differently, there are cases where the same agent has different levels of importance or is not considered important at all. We show the consistency analysis in Table 2.

**Captioning.** Our proposed dataset emphasizes explainability as another aspect. In addition to identifying and ranking important agents, we aim to provide an explanation for

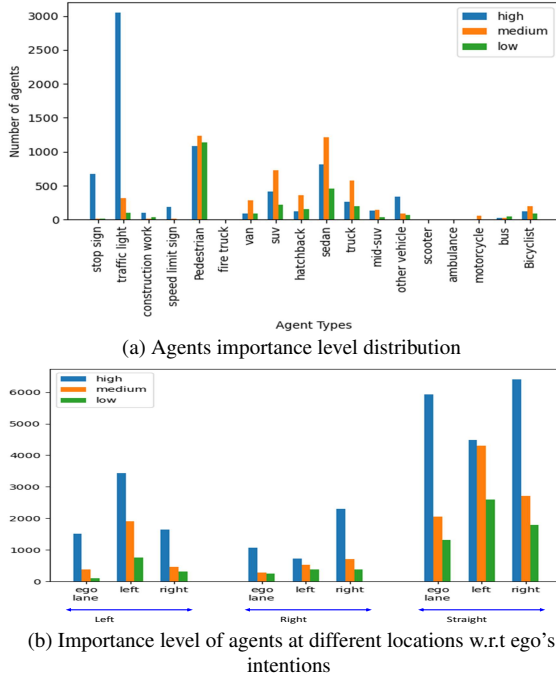


Figure 3. Statistical Analysis of Rank2Tell Dataset

why these objects are deemed significant. Therefore, after annotators identify and localize the important objects, we request them to annotate certain object-level attributes and utilize them to elaborate on why they regard the object as important via a free-form caption. Specifically, as demonstrated in Fig. 2, we request them to annotate the following:

- *What*: What class (type and importance level) does the important agent belong to?
- *Which*: Which visual and motion attribute belongs to the agent?
- *Where*: Where is the agent (location + direction)?
- *How*: How does the ego car respond to the agent?
- *Why*: Why is the agent of High/Medium/Low level of importance?

The first four questions require single-choice answers from a pre-defined set of options, while the last question is an open-ended caption that combines the answers to the previous four questions (what, which, where, how). This allows annotators to use free-form captions while incorporating all essential information captured in the 3W+1H format.

A unique benefit of having multiple captions for a single object is the ability to evaluate caption diversity. Although caption diversity has been explored in image datasets in previous works [8, 33, 38], it has rarely been investigated in a video setting, particularly in traffic scenes. These earlier studies concentrated on measuring semantic diversity by various concepts in the same image. In contrast, we focus on diversity based on how humans perceive importance. Due to the subjective nature of this task, it is crucial to assess how different individuals explain and perceive impor-

tance. As far as we know, our dataset is the first to provide *diverse* captions for *multiple* objects in traffic scenes.

### 3.2. Features and Applications of Rank2Tell

**Scene Graphs.** Our proposed dataset introduces several attributes that offer significant potential for generating informative scene graphs by leveraging spatial (i.e., which and where), temporal (i.e., tracking with 2D+3D point cloud features), and semantic (i.e., what and how) features. These relational attributes are captured by scene graphs, using comprehensive scene information recorded by road users. They are valuable for performing downstream tasks.

**Situational Awareness.** Enhancing situational awareness is critical for safe and efficient navigation in complex traffic scenarios. In a recent study on A-SASS, Wu et al. [39] demonstrated the effectiveness of highlighting important agents in the scene to improve drivers' situational awareness. Motivated by this study, our proposed dataset can potentially be used to identify important agents in the scene and develop an adaptive user interface for improving the driver's situational awareness in real time. Additionally, our dataset offers unique attributes such as *attention* and *communicative aspects* which can be beneficial in Advanced Driver Assistance Systems (ADAS) applications.

**Interpretable Models.** Human-interpretable models are important for safety-critical applications like Autonomous Driving or Advanced Driver Assistance Systems (ADAS). The model's interpretability can be evaluated using comprehensive attributes of *Rank2Tell*, and can potentially be used to address several tasks towards providing explanations of driving risks associated with important agents. Some of these tasks are a) *important agent localization and tracking*, b) *importance level ranking*, c) *caption generation*, and d) *diverse captions generation*. Our dataset also enables joint handling of these tasks, which is another unique aspect. We discuss details of these tasks in supplementary materials.

In this paper, we benchmark the performance on two of these tasks: important agents ranking and captions generation, and also provide a model that jointly addresses them, as discussed in Section 5.

### 3.3. Dataset Analysis

#### 3.3.1 Dataset Statistics

Figure 4 shows the distribution of labels obtained using video-level question answering in Rank2Tell. Figure 3a shows the distribution of agent types with their importance levels in the scene. This is answered using the *what* in the question answering. Since the dataset comprises all scenarios focused on four-way intersections, the majority of the infrastructures such as traffic lights (3048), and stop signs (668) are of high importance and there is consensus among annotators. Figure 3b demonstrates the distribution of location level-1 of various important levels objects given the




Figure 4. Statistical analysis of Rank2Tell dataset.

intention of the ego car. It shows that the majority of agents situated on the left and right of the ego car are of high importance when the ego car’s intention is left and right, respectively. However, when the ego car’s intention is to go straight, the agents’ location is equally distributed in three lanes- left, right, and ego lane. This makes the task of importance ranking classification nontrivial and difficult to estimate by only using the ego car’s intention as a feature.

To describe the visual attributes of significant objects, annotators usually provide free-form responses that integrate information about the object’s what, where, how, and which. In Figure 4, we illustrate the distribution of the top 30 words utilized in captioning the “why” question. This demonstrates that annotators effectively conveyed the intention, motion direction, and location of important agents while generating natural language captions.

### 3.3.2 Consistency Analysis

We conduct an inter-annotator consistency analysis among five annotators, based on the mode of various importance levels they selected, as presented in Table 2. In cases with multiple modes, we set the highest importance level as the final importance to obtain these consistency scores. A 40%, 60%, 80%, and 100.00% consistency implies that 2, 3, 4, and 5 annotators provided the same importance levels for an object. Results show that for data instances with *High* as the majority importance, 88.23% of data samples exhibit more than 60% consistency. Similarly, for data instances with *Not-Important* as the majority voting, 98.61% data samples have more than 60% consistency. This indicates that annotators highly agreed on objects selected as *High* and *Not-Important* importance based on mode. To assess the quality of the dataset, we compute intra-class correlation (ICC) [34] for the entire dataset, yielding 0.92, which shows excellent inter-rater agreement [7]. However, we use the mode of only importance level classes to obtain the ground-truth importance level of an object for various tasks, such as object importance level classification (Section 5.1). That is, if 2 out of 5 annotators deem an agent as important, we use the ground truth as the mode of the two importance levels instead of all 5. This method aims to reduce falsely underes-

timating an agent’s importance. Please refer to the supplementary materials for more details.

Table 2. Data (%) for different inter-annotator consistency (%) of agent’s importance annotation based on the mode of all 5 annotators’ importance levels.

Consistency	Not Important	Low	Medium	High
40	1.37	38.33	33.77	11.75
60	4.52	51.54	42.39	30.93
80	7.86	9.46	17.79	28.95
100	86.23	0.65	6.03	28.35
≥60	98.61	61.65	66.21	88.23

## 4. Methodology

We introduce a model to jointly address objects’ importance classification and caption generation. The architecture of the model is illustrated in Figure 5, which leverages multiple modalities (i.e., 2D and 3D features). It consists of five main components: a) a 2D deep feature extraction module that extracts object-level features from frontal-view visual observations and ego vehicle state information, b) a 3D deep feature extractor that extracts object-level features from 3D point cloud observations, c) a relational graph module that learns the enhanced objects (node) features and relational (edge) features, d) an importance classification module that takes in the objects and objects relational features to predict the importance level for each important objects in the scene, and e) a context-aware attention captioning module that generates descriptions from the object features and relational features.

### 4.1. Feature Extractor

We adopt a 2D deep feature extractor from [22]. The deep feature extractor uses a sequence of RGB images, depth images, semantic maps, the ego vehicle’s inertial data, and a 2D bounding box of each object in the scene. The depth images are obtained by projecting the point cloud to the stitched camera view and the segmentation maps are obtained using DeepLabv3 on the stitched RGB images. To extract the 3D features of each object, we adapt the Pointnet++ [31] backbone and the voting module in VoteNet [11]




Figure 5. **Architecture of Joint Model.** The framework consists of five components: a 2D feature extractor, a 3D feature extractor, a relational graph, an importance classifier, and a captioning decoder.

to aggregate all object candidates in a scene to individual clusters. To capture the mutual influence and relations among objects, we use a graph-based approach that models objects as nodes and their relations as edges in a relational graph module. The module takes in the concatenated 2D and 3D object features and extracts both object features and relational features between objects. To model the relational (edge) features, the module considers only the  $K$  nearest objects surrounding the target object to limit the computation complexity. The final object features are obtained by concatenating the graph node features with the object relations, global features, ego features, and ego intentions. These are then fed into the importance classifier and the captioning decoder, respectively. More details can be found in the supplementary materials.

## 4.2. Training Objective

The loss function for our joint model consists of importance classification loss and caption generation loss. We apply a conventional cross-entropy loss on the generated token probabilities, as in previous works [6, 22]. We used a weighted sum of both loss terms as the final loss:

$$\mathcal{L} = \alpha \mathcal{L}_{\text{imp}} + \beta \mathcal{L}_{\text{cap}}, \quad (1)$$

where  $\alpha$  and  $\beta$  are the weights for the loss terms. Further, to enforce the model to reduce the instances of falsely underestimating an agent’s importance, we penalize the  $\mathcal{L}_{\text{imp}}$  corresponding to different ground truth (GT) and predictions (P) for different importance levels of high (H), medium (M), low (L), and not-importance (NI) as follows:

$$\mathcal{L}_{\text{imp}} = \sum_{i=1}^N \mathcal{L}_i \quad (2)$$

$$\mathcal{L}_i = \begin{cases} \lambda_k \mathcal{L}_i, & \text{if } (GT - P) = k > 0, \\ \mathcal{L}_i, & \text{otherwise,} \end{cases} \quad (3)$$

where  $\mathcal{L}_i$  is the cross-entropy loss for each object  $i$ .

Method	F1 (I)	F1 (NI)	Accuracy
OIE [14]	55.78	87.74	80.80
INTERACT [47]	56.42	88.02	81.21
IOI [22]	64.01	89.06	83.22
Ours	<b>78.44</b>	<b>92.97</b>	<b>89.39</b>

Table 3. Quantitative evaluation of F1 scores for two importance levels across baselines (I: important, NI: nonimportant).

## 5. Experiments

We evaluate the performance of our model on the proposed Rank2Tell dataset by comparing it with various baselines on two tasks: importance level ranking (classification) and natural language captions generation. For fair comparisons, we use the same backbone modules and hyperparameters across different baselines. More details can be found in the supplementary materials.

### 5.1. Importance Level Classification: Baselines

All the importance classification baselines take 2D image features and object features as inputs and infer the importance level of objects in the scene.

**Goal-oriented object importance estimation (OIE) [14]:** OIE is a two-stage framework that firstly generates object tracklets from videos as object proposals and then classifies the proposals as important objects.

**Interaction graphs based object importance estimation (INTERACT) [47]:** INTERACT is a graph convolutional network based method that extracts appearance features from objects and represents them as nodes in the graph. These features are then updated by interacting with other nodes through graph convolution based on the learned interaction graph edges. The updated features are used to estimate object importance using a multi-layer perceptron.

**Important Object Identification (IOI) [22]:** IOI is a graph convolutional network based method that explicitly models both the appearance and motion of objects in the scene and also uses the ego vehicle state information in the

Method	F1 (L)	F1 (M)	F1 (H)	F1 (NI)	Accuracy
OIE [14]	10.25	14.32	49.68	86.56	74.95
INTERACT [47]	14.04	27.56	44.04	87.45	75.09
IOI [22]	14.09	16.19	49.55	87.85	76.84
Ours	<b>20.49</b>	<b>28.94</b>	<b>58.84</b>	<b>92.62</b>	<b>80.93</b>

Table 4. Quantitative evaluation comparing the F1 scores for 4 importance levels across baselines. L:LOW, M:MEDIUM, H:HIGH, NI:NON-IMPORTANT

Method	C	B-4	M	R
S&T [37]	47.67	30.24	34.91	53.31
Scan2Cap [6]	56.32	<b>49.59</b>	<b>38.36</b>	66.35
Ours	<b>100.15</b>	45.83	36.21	<b>68.56</b>

Table 5. Quantitative evaluation comparing the performance of baselines for captions predictions. C: CIDER, B-4: Bleu-4, M: Meteor, R: Rouge

current frame for importance classification. It reasons about the relations between objects in the scene.

## 5.2. Captioning: Baselines

**Scan2Cap [6]:** Scan2Cap is an end-to-end method to perform dense captioning on 3D point clouds to densely detect and describe 3D objects in RGB-D scans. The method first employs a detection pipeline to obtain object proposals and then applies a relational graph and context-aware attention captioning module to learn object relations and generate tokens, respectively.

**Show, and Tell (S&T) [37]:** This baseline generates captions using 2D features (i.e., global RGB image features and target object features). The visual features are extracted using a ResNet-101 pre-trained on the ImageNet dataset. The global features are concatenated with the target object features, which are used to generate the captions.

## 5.3. Metrics

To measure the quality of the predicted importance ranking of each object in the scene and their corresponding captions, we evaluate the performance of the baselines using standard classification metrics: F1 score, and accuracy for the importance estimation, and standard metrics such as BLEU-4 (B4) [30], METEOR (M) [2], ROGUE (R) [26], and CIDER (C) [36].

## 5.4. Results

### 5.4.1 Quantitative Comparison

Table 3, Table 4, and Table 5 show the quantitative results of our model and baselines for importance classification and caption prediction tasks. The importance baselines are trained and evaluated with two classes (Important and Non-Important), and four classes (High, Medium, Low, and

Non-Important). For the two classes, high, medium, and low important objects are merged as one important class. The INTERACT [47] baseline outperforms OIE [14] as it models agent interactions in the scene using graph convolutional networks. This makes it more efficient in predicting the importance levels of multiple agents in the scene with lesser input features than OIE [14]. Furthermore, IOI [22] outperforms the previous two baselines as it utilizes the ego vehicle’s speed, acceleration, and intention features. While the agent’s importance level is highly influenced by the ego vehicle’s intention (e.g., straight, left, right) at the intersection, the ego’s speed and acceleration are influenced by the importance level of an agent. This interaction behavior is explicitly incorporated in IOI [22], which makes it outperform other methods. Since there is currently no importance ranking estimation model that uses 3D features, we benchmark results on baselines using only 2D features.

For the captioning baselines, Scan2cap [6] outperforms S&T [37] because Scan2Cap uses 3D point clouds and 3D bounding box features, which are more informative than 2D features. The 3D point cloud provides information on the distance between agents and their heights, which captures the dependence of the agents’ location with the words in the captions, such as *left side* or *right side* of the ego car. Our proposed joint model leverages both 2D and 3D features and outperforms the uni-modal baselines in several metrics. Further, the joint training of importance classification and captions prediction complements each other, thereby outperforming the baselines trained with a single task (more details in supplementary material).

To demonstrate the usefulness of Rank2Tell annotations, we conduct an ablation study where we integrate action attributes (*Which*) with object features for all baselines, and the detailed results are shown in supplementary materials.

## 6. Conclusion

We present a novel multi-modal dataset (Rank2Tell) in urban traffic scenarios, which enables joint importance level prediction and reasoning in traffic scenes with frame-level annotations. The dataset includes object-level questions on important objects and language interpretations of the scene from the ego driver’s perspective, integrating spatial, temporal, and relational features. This offers new avenues for improving visual scene comprehension and advancing autonomous systems’ interpretability and trustworthiness.

## References

- [1] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6077–6086, 2018. 2



- [2] 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. 8
- [3] Keivan Borna and Reza Ghanbari. Hierarchical lstm network for text classification. *SN Applied Sciences*, 1:1–4, 2019. 2
- [4] 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. 3
- [5] Dave Zhenyu Chen, Angel X Chang, and Matthias Nießner. Scanrefer: 3d object localization in rgb-d scans using natural language. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX*, pages 202–221. Springer, 2020. 2
- [6] Zhenyu Chen, Ali Gholami, Matthias Nießner, and Angel X Chang. Scan2cap: Context-aware dense captioning in rgb-d scans. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3193–3203, 2021. 7, 8
- [7] Domenic V Cicchetti. Guidelines, criteria, and rules of thumb for evaluating normed and standardized assessment instruments in psychology. *Psychological assessment*, 6(4):284, 1994. 6
- [8] Bo Dai, Sanja Fidler, Raquel Urtasun, and Dahua Lin. Towards diverse and natural image descriptions via a conditional gan. In *Proceedings of the IEEE international conference on computer vision*, pages 2970–2979, 2017. 5
- [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. 3
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2
- [11] Zhipeng Ding, Xu Han, and Marc Niethammer. Votenet: A deep learning label fusion method for multi-atlas segmentation. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part III 22*, pages 202–210. Springer, 2019. 6
- [12] Partners for Automated Vehicle Education. Pave poll 2020, 2020. 1
- [13] Mingfei Gao, Ashish Tawari, and Sujitha Martin. Goal-oriented object importance estimation in on-road driving videos. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 5509–5515. IEEE, 2019. 2
- [14] Mingfei Gao, Ashish Tawari, and Sujitha Martin. Goal-oriented object importance estimation in on-road driving videos. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 5509–5515. IEEE, 2019. 7, 8
- [15] Zhong Ji, Yuxiao Zhao, Yanwei Pang, Xi Li, and Jungong Han. Deep attentive video summarization with distribution consistency learning. *IEEE transactions on neural networks and learning systems*, 32(4):1765–1775, 2020. 2
- [16] Justin Johnson, Andrej Karpathy, and Li Fei-Fei. Denscap: Fully convolutional localization networks for dense captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4565–4574, 2016. 2
- [17] Jinkyu Kim and John Canny. Interpretable learning for self-driving cars by visualizing causal attention. In *Proceedings of the IEEE international conference on computer vision*, pages 2942–2950, 2017. 2
- [18] Jinkyu Kim and John Canny. Interpretable Learning for Self-driving Cars by Visualizing Causal Attention. In *ICCV*, 2017. 2, 3
- [19] Jinkyu Kim, Anna Rohrbach, Trevor Darrell, John Canny, and Zeynep Akata. Textual explanations for self-driving vehicles. In *Proceedings of the European conference on computer vision (ECCV)*, pages 563–578, 2018. 2, 3
- [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*, 123:32–73, 2017. 2
- [21] Chengxi Li, Stanley H Chan, and Yi-Ting Chen. 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. 2
- [22] Jiachen Li, Haiming Gang, Hengbo Ma, Masayoshi Tomizuka, and Chiho Choi. Important object identification with semi-supervised learning for autonomous driving. In *ICRA*, 2022. 1, 2, 4, 6, 7, 8
- [23] Jiachen Li, Fan Yang, Hengbo Ma, Srikanth Malla, Masayoshi Tomizuka, and Chiho Choi. Rain: Reinforced hybrid attention inference network for motion forecasting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16096–16106, 2021. 1
- [24] Qiang Li, Chunsheng Liu, Faliang Chang, Shuang Li, Hui Liu, and Zehao Liu. Adaptive short-temporal induced aware fusion network for predicting attention regions like a driver. *IEEE Transactions on Intelligent Transportation Systems*, 23(10):18695–18706, 2022. 2
- [25] Wei-Hong Li, Fa-Ting Hong, and Wei-Shi Zheng. Learning to learn relation for important people detection in still images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5003–5011, 2019. 2
- [26] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004. 8
- [27] Xiaobai Ma, Jiachen Li, Mykel J Kochenderfer, David Isele, and Kikuo Fujimura. Reinforcement learning for autonomous driving with latent state inference and spatial-temporal relationships. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6064–6071. IEEE, 2021. 1
- [28] 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 Confer-*

- ence on Applications of Computer Vision, pages 1043–1052, 2023. 2, 3, 4
- [29] Eshed Ohn-Bar and Mohan Manubhai Trivedi. Are all objects equal? deep spatio-temporal importance prediction in driving videos. *Pattern Recognition*, 64:425–436, 2017. 2, 3, 4
- [30] 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. 8
- [31] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017. 6
- [32] Vasili Ramanishka, Yi-Ting Chen, Teruhisa Misu, and Kate Saenko. Toward driving scene understanding: A dataset for learning driver behavior and causal reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7699–7707, 2018. 2, 3
- [33] Rakshith Shetty, Marcus Rohrbach, Lisa Anne Hendricks, Mario Fritz, and Bernt Schiele. Speaking the same language: Matching machine to human captions by adversarial training. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4135–4144, 2017. 5
- [34] Patrick E Shrout and Joseph L Fleiss. Intraclass correlations: uses in assessing rater reliability. *Psychological bulletin*, 86(2):420, 1979. 6
- [35] Guolei Sun, Wenguan Wang, Jifeng Dai, and Luc Van Gool. Mining cross-image semantics for weakly supervised semantic segmentation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 347–365. Springer, 2020. 2
- [36] 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. 8
- [37] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3156–3164, 2015. 8
- [38] Qingzhong Wang and Antoni B Chan. Describing like humans: on diversity in image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4195–4203, 2019. 5
- [39] Tong Wu, Enna Sachdeva, Kumar Akash, Xingwei Wu, Teruhisa Misu, and Jorge Ortiz. Toward an adaptive situational awareness support system for urban driving. In *2022 IEEE Intelligent Vehicles Symposium (IV)*, pages 1073–1080. IEEE, 2022. 2, 3, 4, 5
- [40] 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. 2
- [41] Yiran Xu, Xiaoyin Yang, Lihang Gong, Hsuan-Chu Lin, Tz-Ying Wu, Yunsheng Li, and Nuno Vasconcelos. Explainable object-induced action decision for autonomous vehicles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9523–9532, 2020. 2, 3
- [42] 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. 2
- [43] 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. 3
- [44] 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. 2
- [45] Zhou Yu, Jun Yu, Yuhao Cui, Dacheng Tao, and Qi Tian. Deep modular co-attention networks for visual question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6281–6290, 2019. 2
- [46] 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. 3
- [47] Zehua Zhang, Ashish Tawari, Sujitha Martin, and David Crandall. Interaction graphs for object importance estimation in on-road driving videos. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8920–8927. IEEE, 2020. 1, 7, 8
- [48] Zehua Zhang, Ashish Tawari, Sujitha Martin, and David Crandall. Interaction graphs for object importance estimation in on-road driving videos. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8920–8927. IEEE, 2020. 2