

## A. Implementation Details

We present additional implementation details to reproduce our baseline experiments below. We primarily consider two-stage baselines that first track all objects and then classify them into `referred` objects, `related` objects, or `other` objects. We repurpose 3D trackers from winning teams of the Argoverse 2 End-to-End Forecasting Challenge. All models are trained to output tracks for 26 annotated categories within Argoverse 2. We include an example prediction signature below:

```
1 {
2     "timestamp_ns": int,
3     "track_id": int,
4     "confidence": float,
5     "class_name": str,
6     "translation_m": np.array,
7     "size": np.array,
8     "yaw": np.array,
9 }
```

We append a track for the ego vehicle with a consistent `track_id`, confidence of 1.00, class name of `EGO_VEHICLE`, yaw of 0 radians, size (`<length, width height>` in meters) of `<4.877, 2, 1.473>`, and translation (`<x, y, z>` in meters) of `<1.422, 0, 0.25>`. This translation represents the offset from the centroid of the ego-vehicle to the ego-vehicle reference coordinate near the rear axle.

**All Tracks Oracle.** Our naive baseline labels all ground truth tracks as `referred` objects. This means that all logs and timestamps are labeled as “positive”, giving a balanced accuracy of 50%.

**Filter by Referred Class.** This baseline filters all tracks except those corresponding to the referred class. To get the referred class, we prompt Claude 3.7 Sonnet as follows:

```
Categories: {AV2 Categories}
Please select the one category that most corresponds to the
object of focus in the description:{Natural Language De-
scription}. As an example, for the description “vehicle
turning at intersection with nearby pedestrians” your output
would be VEHICLE. For the description “ego vehicle near
construction barrel” your output would be EGO_VEHICLE.
Your only output should be the category name.
```

We define two super-categories for `VEHICLE` and `ANY_OBJECT`. The `VEHICLE` category is manually defined to contain the classes `ARTICULATED_BUS`, `BOX_TRUCK`, `BUS`, `EGO_VEHICLE`, `LARGE_VEHICLE`, `MOTORCYCLE`, `RAILED_VEHICLE`, `REGULAR_VEHICLE`, `SCHOOL BUS`, `TRUCK`, and `TRUCK_CAB`.

**Image Embedding Similarity.** This baseline filters tracks by computing prompt-track CLIP similarity over time. At each timestep, we project the eight vertices of each ob-

ject’s bounding box onto the image plane of all seven cameras. We then cull any objects that do not partially fall within the frustum of each camera. For each projected bounding box, if three or more vertices fall within the image plane, we take a crop using the minimum and maximum uv image coordinates. We pad this crop by 30 pixels on each side and feed the resulting image into a CLIP ViT-L/14 model to obtain CLIP image embeddings. We also use CLIP ViT-L/14 to compute text embeddings for all prompts in the dataset. Next, we calculate cosine similarity scores between each track’s image crops and the text prompt. Since individual similarity scores tend to be noisy, we follow ReferGPT [6] in applying a dynamic threshold and majority voting strategy to identify the `referred` object tracks. Since the dataset includes both positive and negative examples for each prompt, we group similarity scores by prompt and threshold the top 10% of scores. If more than half of the scores of a track fall within the top 10% threshold, we label the entire track as a `referred` object. To reduce the number tracks of length one (i.e., tracks spanning only one timestamp), we apply a similarity score modifier:  $s_{\text{modified}} = s_{\text{cosine}} - \frac{0.05}{\text{len}(\text{track})}$ .

**LLM APIs as a Black Box** We use the [OpenAI Responses API](#) to determine whether particular scenarios occur within a data log. We prompt GPT-5 with three files: an HD map data, 6-DoF ego-to-city transformations, and object tracker predictions. We use `gpt-5-2025-08-07` with medium reasoning effort and access to a code interpreter tool. We give instructions to save a CSV containing the `referred` object id’s and timestamps for each scenario. All 150 test data logs are run in parallel, with one request and one response for each data log. The saved CSVs are parsed into the RefAV format for evaluation as each response is delivered. We include the instructions given to GPT-5 below.

```
You are an expert data analyst. Your job is to identify if a
given scenarios occurs within an autonomous driving log.
A log contains files with an HD map, 6-DoF poses of the
ego vehicle, and bounding box annotations from an object
tracker.
For each of the given scenario descriptions, output a CSV file
containing the track_uuids and timestamps of the referred
objects. Some logs may not contain any referred objects.
Your csv file should have “track_uuid” and “timestamp_ns”
columns. Your output csv files should have the naming
convention scenario{i}.csv, depending on the order the de-
scriptions are given. Make sure to save all of your CSVs!
Identify if each of these scenarios occurs within the log.
{scenario_descriptions}
```

**ReferGPT.** This baseline uses a VLM to caption the cropped images of each track at each timestep. Tracks are filtered by computing a text similarity score between the prompt and track captions. We take several steps to adapt

ReferGPT to the RefAV dataset. At each timestep, we project the eight vertices of each predicted track’s bounding box onto the image plane of all seven cameras. We then cull any objects that do not partially fall within the frustum of each camera. For each projected bounding box, if three or more vertices fall within the image plane, we take a crop using the minimum and maximum uv image coordinates. If the predicted track is visible in multiple cameras, we use the 2D bounding box with the highest area. We pad this crop by 10% of the crop’s length on each side. For the ego-vehicle, the crop is the entire image from the front camera. Each of the per object-timestamp crops is also paired with the motion statistics. We compute a track  $i$ ’s motion statistics at timestep  $t$  as  $\mathbf{C}_t^i = [x, y, \theta, v_x, v_y, \alpha, d]$ , where  $x$  and  $y$  are the position of the track in meters,  $\theta$  is the yaw of the track in degrees,  $v_x$  and  $v_y$  are the velocities of the track in meters per second,  $\alpha$  is the yaw rate of the track in degrees per second.  $d$  is the distance from the centroids of the ego-vehicle and track in meters. All values are with respect to a static ego-centric coordinate frame. We prompt gpt-5-mini to caption each of the images as follows:

Caption this image taken from an autonomous vehicle (referred to as the ego-vehicle). Our detection model classifies this object as {object\_category}.  
 Here are the coordinates of the object {object\_motion}. This is in the format [x, y, yaw, x\_velocity, y\_velocity, angular\_velocity, distance].  
 This is in a coordinate system where positive x is to the front of the ego vehicle, positive y is to the left of the ego-vehicle, and positive z is above the ego-vehicle. Position and velocity are measured in meters and meters per second. Yaw is measured in degrees counter-clockwise (colloquially “left”) from the x-axis. Angular velocity is in degrees per second. Distance is distance from the ego-vehicle.  
 Please describe the appearance of the object and scene, if they are located relative to the road/ego-vehicle, how they are moving, and how they are interacting with any nearby objects (if any are visible). Treat values of x, y, x\_velocity, and y\_velocity near zero as zero. Feel free to add more relevant information in the captions.

We present an example caption below:

A motorcyclist slightly to the front-right of the ego-vehicle is stopped or creeping forward beside a compact SUV, straddling a scooter with feet down and hands on the handlebars. The rider is close to the SUV’s rear quarter, appearing stationary or moving very slowly while sharing the lane alongside other stopped traffic. The scene looks like an urban intersection or congested street with vehicles side-by-side and the rider positioned between them.

To determine if an object is referred, we compute a score  $s = s_{\text{cosine}} + 0.1s_{\text{fuzzy}}$ .  $s_{\text{cosine}}$  is the cosine similarity between siglip2-base-patch16-naflex text embeddings of the caption and prompt. We follow the original ReferGPT implementation and compute  $s_{\text{fuzzy}}$  using Gestalt pattern matching to find the similarity between two strings. Finally, we determine that an entire track is referred if the majority the track’s caption scores are above 0.63.

We find that the SigLIPv2 text similarity does not effectively separate the referred and related objects in the prompt. For example, the query “car to the left of a pedestrian”, has high similarity with both “the vehicle is to the left of the person” and “the person is to the right of the vehicle”. To resolve this issue, we post-process predictions with *Filter by Referred Class* to obtain the final set of referred objects.

**Referential Tracking by Program Synthesis.** RefProg filters off-the-shelf tracks using Python programs synthesized by a large language model (LLM), without explicitly relying on any information from LiDAR or cameras. To reduce the time required to run these synthesized programs, we first filter the input tracks based on confidence. We compute each track’s confidence score by summing its confidence values across all timestamps. For the classes REGULAR\_VEHICLE, PEDESTRIAN, BOLLARD, CONSTRUCTION\_CONE, and CONSTRUCTION\_BARREL, we retain the top 200 tracks per class. For all other classes, we retain the top 100 tracks per category. We prompt the LLM as follows:

Please use the following functions to find instances of a referred object in an autonomous driving dataset. Be precise to the description, try to avoid returning false positives.  
 API Listing: {API Listing}  
 Categories: {AV2 Categories}  
 Define a single scenario for the description:{Natural Language Description}  
 Here is a list of examples: {Prediction Examples}. Only output code and comments as part of a Python block. Do not define any additional functions, or filepaths. Do not include imports. Assume the log\_dir and output\_dir variables are given. Wrap all code in one python block and do not provide alternatives. Output code even if the given functions are not expressive enough to find the scenario.

We include the API listing in Appendix J. We found this prompt consistently produces reasonable code across all tested LLMs. After running the synthesized programs, we post-process the output tracks. We discard object relationships where the spatial distance exceeds 50 meters. For the remaining tracks, we apply timestamp dilation to reduce flickering of the referred object labels. We symmetrically dilate time segments to a minimum of 1.5 seconds.

For example, if the base tracker tracks an object with ID 1 from 3.0 to 15.0 seconds at 2 Hz, and the synthesized code marks it as referred at timestamps 4.5 and 5.0, we update the final output to label the track as referred from 4.0 to 5.5 seconds. Finally, we downsample all output tracks to 2 Hz for evaluation.

## B. Token Statistics for LLM Baselines

*ReferGPT*, *LLM APIs as a Black Box*, and *RefProg* can be viewed as different ways of providing scaffolding around an LLM for RMOT and scenario mining tasks. Most commercially available LLM APIs bill according to four cost categories: uncached input tokens, cached input tokens, output tokens, and code interpreter sessions (CIS). Input token count scales approximately linearly with the length of the input text, the number of input images, and seconds of input video. *ReferGPT* requires orders of magnitude more input and output tokens because it calls a VLM call for every tracked object at every timestamp. It also requires processing images, which can use more than 1000 tokens (depending on resolution) using GPT-5-mini. *LLM APIs as a Black Box* only makes as many calls as are there logs of data. However, inspecting the reasoning traces reveals that GPT-5 automatically went through multiple rounds of generating data processing code and verifying the output. The LLM also often re-generates the same processing code between logs. Initializing code interpreter sessions represents approximately half of the total cost (\$38.49) of inference. *RefProg* makes as many calls as there are unique scenario descriptions. Since the entire atomic function API must be given in context, it requires a relatively high number of input tokens. However, most of these tokens are cached between requests. The output of the *RefProg* is relatively compact because the LLM has access to high-level composable functions and does not need to write extensive Python code.

Table 5. **Token Efficiency by LLM Baseline** *RefProg* achieves the highest scenario mining performance for the lowest cost. Cost is in USD for the OpenAI Platform in November 2025. We report the total number of tokens used millions. CIS refers to code interpreter sessions.

Method	LLM	HOTA-Temporal ↑	Cost ↓	Uncached Input	Cached Input	Output	CIS
<i>ReferGPT</i>	GPT-5-mini	20.7	150.64	372.9M	4.0M	28.7M	0
LLM API as Black Box	GPT-5	37.2	77.59	4.5M	0	3.4M	1283
<i>RefProg</i> (Ours)	GPT-5	42.3	13.89	0.6M	2.6M	1.3M	0

## C. Comparison of Recent Datasets

We provide a qualitative comparison of language expressions found in *RefAV* with other popular RMOT-for-driving and VQA-for-driving datasets. In particular, we examine *nuPrompt*, *Refer-KITTI*, and *nuScenes-QA*. To obtain a representative sample of each dataset, we compute text embeddings for each expression using Qwen-3-Embedding-8B. We

display the first 30 samples using farthest point sampling initialized from the two samples furthest apart from each other. Many expressions in *nuPrompt* deal with object attributes that are minimally relevant to driving such as car color, a person’s gender, or a person’s clothing. *Refer-KITTI* similarly only refers to annotated vehicle and pedestrian classes. *nuScenes-QA* asks questions that are relevant to driving, but the dataset is limited by its lack of human annotations and unnatural sounding questions. We present examples from each dataset below.

### RefAV (Ours)

- the vehicle two cars behind the ego vehicle
- pedestrian walking with dog on sidewalk
- stationary object
- motorcycle at stop sign
- group of people
- construction workers working next to the road
- fire truck in a non-emergency
- excavator
- large truck blocking view of vehicle that is about to turn left onto a main road
- bicyclist changing lanes to the right
- car with a stroller to its left
- wheelchair user at pedestrian crossing
- person directing traffic in a school zone
- skateboarder on the road
- unattended dog
- accelerating wheeled devices
- construction barrel with at least 2 construction cones within 3 meters
- bus with multiple pedestrians waiting on the right side
- vehicle passing the ego vehicle going significantly over the speed limit on an urban road
- van facing the wrong way in the middle turn lane
- two of the same work truck
- pedestrian near message board trailer
- pedestrians over 30 meters away
- the vehicle behind another vehicle being crossed by a jaywalking pedestrian
- vehicle flashing their hazard lights
- ego vehicle driving while it is cloudy
- motorcycle or pedestrian within 5 meters to the right of the ego vehicle
- bicycle behind a bus
- scooter lying down on the side of the road
- vehicles being passed by a motorcycle on either side

### nuPrompt [72]

- Back, the pedestrian
- a stationary orange parking is on the parking space
- at the parking space, the trailer
- Car waiting with the same direction
- A person on foot by the roadside
- The silver lightcolored parked car is stationary in the parking space
- The objects possess the color blue
- Males with a black bag are getting ready to cross the road
- The one riding the black lefthandturn bicycle is a person
- Men stand on the street wearing pants and hats
- Females walking across the street
- Men wear pants while walking across the street
- In a state of sitting
- The car is in the act of turning
- Pedestrian getting ready to cross over
- is in the process of making a turn
- The car in motion is black and moving in the opposite direction
- Males wear tshirts and pants while making a left turn on the opposite side of the road
- the pedestrian is situated at the back left
- The dark car is at thestandstill and awaiting with the opposite direction as the truck
- The truck is thepale silver color
- we see the car ahead
- Females walk while wearing pants in the same direction
- Security camera
- Men are wearing shorts while standing on the road
- There is the truck positioned in the parking spot
- Color is red
- The bus is in front
- On the pedestrian pathway
- The car that is moving in the same direction is blue

### Refer-KITTI [71]

- cars which are faster than ours
- walking pedestrian in the left
- women back to the camera
- males in the right
- left cars which are parking
- people in the pants

- people
- women carrying a bag
- light-color vehicles
- vehicles in the counter direction
- walking males
- standing females
- men in the left
- right cars in silver
- red moving cars
- women
- walking people in the right
- vehicles
- pedestrian
- vehicles in front of ours
- vehicles which are in the left and turning
- vehicles in white
- cars in horizon direction
- black vehicles in right
- standing men
- vehicles which are braking
- walking women
- men
- right cars in the same direction of ours
- black cars

### nuScenes-QA [51]

- How many without rider motorcycles are there?
- There is a thing that is both to the front of me and the front left of the traffic cone; what is it?
- Does the trailer have the same status as the truck that is to the back of the trailer?
- How many pedestrians are to the back of the moving bus?
- What is the with rider thing that is to the back right of the with rider bicycle ?
- Are there any things?
- Is there another barrier of the same status as the barrier?
- The car to the back right of the stopped construction vehicle is in what status?
- There is a parked truck; what number of barriers are to the front left of it?
- Are there any other bicycles that in the same status as the car?
- What status is the motorcycle that is to the front left of the standing pedestrian?
- The without rider thing that is to the front of the construction vehicle is what?
- What number of traffic cones are there?

- Are there any trailers to the front of me?
- What is the thing to the back of the stopped thing?
- There is a without rider thing; are there any parked cars to the back of it?
- What number of barriers are to the back of the motorcycle?
- Are there any animals to the back right of me?
- The thing that is to the back right of the truck and the back right of the bus is what?
- There is a bus; is its status the same as the thing that is to the front of the bus?
- There is a with rider bicycle; what number of moving things are to the front of it?
- Are there any moving construction vehicles?
- The standing pedestrian to the back right of the traffic cone is what?
- How many buss are there?
- Is the status of the truck to the back of the pedestrian the same as the truck to the front of the bicycle?
- Is there another thing that has the same status as the animal?
- Are there any things to the back of the moving trailer?
- Are there any moving buss to the front of the without rider thing?
- The thing that is to the front of the moving truck and the front left of the stopped trailer is what?
- How many stopped trailers are there?

## D. Scenario Mining Annotation Tool

In order to manually annotate interesting scenarios, we build a custom web app with Claude 3.7 (Figure 5). To construct a scenario, the user selects an object by clicking on a point within the image frame, writes a natural language description, and selects the start and end frames of the referred objects correspond to the prompt. We project a ray originating from the 2D point and find the ground truth 3D bounding box centroid closest to the ray. This tool allows us to quickly generate multi-object and multi-camera referential tracks.

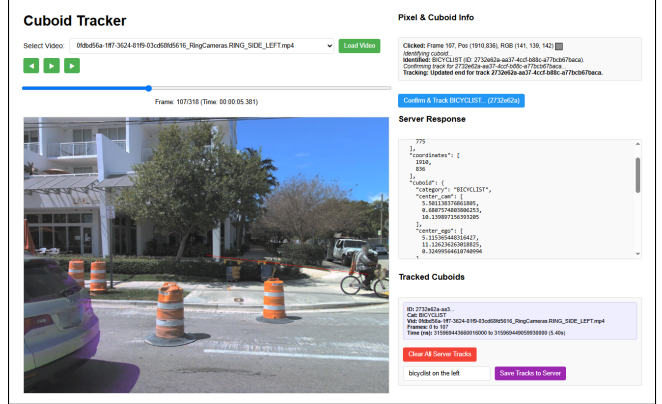


Figure 5. **Manual Annotation Tool.** We create an annotation tool to assist with labeling manually defined scenarios. Our tool allows us to quickly annotate multi-object referential tracks in AV2.

## E. Referential Tracking Evaluation of Related and Other Objects

RefProg creates a scene graph for each natural language prompt that not only identifies referred objects, but also classifies tracks as related objects and other objects. For example, given the prompt vehicle near a pedestrian, the vehicle would be the referred object, the pedestrian would be the related object, and everything else in the scene would be considered other objects. We benchmark the base tracker with HOTA and evaluate HOTA-Temporal of related objects and other objects in Table 6. Since most objects are not related to the prompt, HOTA and HOTA-Temporal Other are similar.

Table 6. **Referential Tracking Accuracy for Related and Other Objects.** We present RefProg’s referential tracking accuracy for related objects and other objects. We find that HOTA (for the base tracker) is similar to HOTA-Temporal (Other) because most objects are not relevant to the referential prompt.

Method	HOTA $\uparrow$	HOTA-Temporal (Related) $\uparrow$	HOTA-Temporal (Other) $\uparrow$
<b>Referential Tracking by Program Synthesis</b>			
Ground Truth	100.0	57.2	97.9
Le3DE2E [8]	78.9	35.6	74.2
ReVoxelDet [28]	69.4	28.2	63.4
TransFusion [2]	69.5	28.6	63.8
Valeo4Cast [74]	75.1	30.8	69.2

## F. Impact of Sampling Rate on RefProg

We evaluate the impact of sampling rate on RefProg. First, we subsample three detectors (Le3DE2E [8], Valeo4Cast [74], and BEVFusion [37]) at 2Hz and 10 Hz. Next, we associate all detections with AB3DMOT [69] to generate 2Hz and 10Hz tracks respectively. Lastly, we evaluate RefProg on all trackers at 2Hz. We find that standard tracking

performance drops by 5% when subsampling due to the difficulty of associating sparse detections. We see a similar drop in HOTA-Track, Log Balanced Accuracy and Timestamp Balanced Accuracy. Interestingly, we do not see a similar drop in performance with HOTA-Temporal. This suggests that RefProg might be robust to temporally sparse inputs and ID switches.

Table 7. **Impact of Sampling Rate.** We evaluate RefProg’s performance with subsampled inputs to measure the relative impact on referential tracking accuracy. We find that standard tracking performance drops by 5% when subsampling due to the difficulty of associating sparse detections.

Method	Hz	HOTA ↑	HOTA-Temporal ↑	HOTA-Track ↑	Log Bal. Acc. ↑	Timestamp Bal. Acc. ↑
Ground Truth	10	100.0	64.7	68.7	81.1	80.7
Ground Truth	2	100.0	64.2	68.8	80.6	80.1
Le3DE2E Det. [8] + AB3DMOT [69]	10	78.0	48.3	49.2	73.1	73.6
Le3DE2E Det. [8] + AB3DMOT [69]	2	73.0	45.2	44.1	69.3	69.5
ValcooCast Det. [17] + AB3DMOT [69]	10	74.8	42.8	45.5	71.6	71.8
ValcooCast Det. [17] + AB3DMOT [69]	2	69.9	45.2	44.1	69.3	69.5
BEVFusion Det. [37] + AB3DMOT [69]	10	76.0	46.4	47.8	73.7	72.3
BEVFusion Det. [37] + AB3DMOT [69]	2	70.9	46.4	45.2	72.1	72.9

## G. More Dataset Statistics

Figure 6 quantifies important characteristics of RefAV such as the number of unique scenario descriptions, the number of positive and negative examples, the number of procedurally and manually defined scenarios, the number of scenarios that include object relationships, and the number of scenarios directly involving the ego-vehicle. We use the catch-all term “expression” for scenario description to better compare RefAV to other datasets. A positive match occurs when the data log contains at least one referred object corresponding to the scenario description. Figure 7 shows the number of positive and negative examples for representative scenario descriptions. In Figure 8, we evaluate the spatial distribution of referred objects.

## H. Summary of Competition Top Performers

We summarize the contributions of top teams below.

**Team Zeekr UMCV** extends RefProg with global context-aware generation to processes all queries collectively for improved reasoning and consistency. They also propose a multi-agent refinement system where a secondary refiner agent iteratively debugs and enhances the generated code.

**Team Mi3 UCM** proposes two key improvements to RefProg: (1) Fault-Tolerant Iterative Code Generation (FT-ICG), which refines faulty code by re-prompting the LLM with error feedback until successful execution, and (2) Enhanced Prompting for Spatial Relational Functions (EP-SRF), which clarifies function semantics to prevent misinterpretations of spatial relationships between objects.

**Team ZXH** refines RefProg by prompting LLMs with a batch of queries and in-context examples instead of prompting for each query one at a time. Importantly, this team’s approach relies on in-context learning and batch prompting to improve consistency and efficiency.

Expression Type	# Expressions
Unique	407
Positive Match	7910
Negative Match	2090
Procedurally Defined	9880
Manually Defined	120
Includes Related Objects	5531
Ego is the only Referred Object	827
Ego is the only Related Object	1200




Figure 6. **RefAV Expression Statistics.** We compute the number of referring expressions categorized by expression type, count the number of the expressions that are positive or negative, and whether they are procedurally generated or manually constructed. We also highlight the number of positive expressions that include related object annotations and explicit references to the ego vehicle. The bar chart on the right shows the number of objects referred to by each expression. Most expressions in RefAV refer to zero (e.g. a negative match) or one object. On average, each expression refers to 3.92 objects.

## I. Examples of RefProg’s Synthesized Programs

We provide examples of RefProg’s correctly synthesized programs below.

*Prompt: Pedestrian crossing between stopped buses*

```

1 # Description: pedestrian crossing between
  stopped buses
2
3 # First, let's get all pedestrians and buses from
  the dataset
4 pedestrians = get_objects_of_category(log_dir,
  category="PEDESTRIAN")
5 buses = get_objects_of_category(log_dir,
  category="BUS")
6
7 # We need to find stopped buses
8 stationary_buses = stationary(buses, log_dir)
9
10 # Alternatively, we could use has_velocity with
  a low threshold to find buses that are
  temporarily stopped
11 stopped_buses = has_velocity(buses, log_dir,
  min_velocity=0, max_velocity=0.5)
12

```

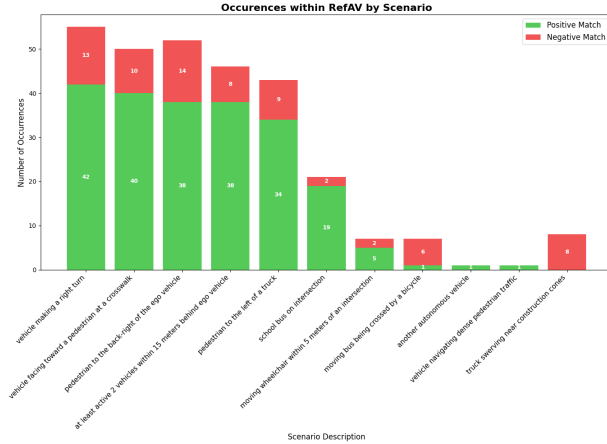


Figure 7. **Occurrences within RefAV by Scenario.** RefAV contains 407 unique scenario descriptions. We display 11 representative descriptions picked uniformly. After the diversity sampling process, RefAV includes about 40 positive matches for common scenarios such as vehicle making a right turn. Less common scenarios such as moving bus being crossed by a bicycle occur only in only a handful of logs across the ArgoVerse 2 dataset. RefAV includes negative examples for almost all procedurally defined scenarios. Manually annotated scenarios such as another autonomous vehicle do not have corresponding negative examples.

```

13 # Now we need to find pedestrians that are
14 # crossing between the stopped buses. First,
15 # identify pedestrians that are crossing in
16 # front of a stopped bus
17 peds_crossing_bus = being_crossed_by(
18     stationary_buses, pedestrians, log_dir)
19 # Reverse the relationship to get the pedestrians
20 # that are doing the crossing
21 crossing_peds = reverse_relationship(
22     being_crossed_by)(stationary_buses,
23     pedestrians, log_dir)
24 # Now we need to find pedestrians that have buses
25 # on both sides (left and right). This means
26 # the pedestrian is between buses
27 peds_with_bus_left =
28     has_objects_in_relative_direction(
29     crossing_peds, stationary_buses, log_dir,
30     direction="left", min_number=1,
31     within_distance=10)
32 peds_with_bus_right =
33     has_objects_in_relative_direction(
34     crossing_peds, stationary_buses, log_dir,
35     direction="right", min_number=1,
36     within_distance=10)
37 # Combine to find pedestrians with buses on both
38 # sides
39 peds_between_buses = scenario_and([
40     peds_with_bus_left, peds_with_bus_right])
41 # Output the scenario
42 output_scenario(peds_between_buses, "pedestrian

```

```

crossing between stopped buses", log_dir,
output_dir)

```

*Prompt: Group of at least 3 moving bicyclists with 5m from each other*

```

1 # Description: Group of at least 3 moving
2 # bicyclists within 5 meters of each other
3 # Get all bicyclists from the dataset
4 bicyclists = get_objects_of_category(log_dir,
5     category="BICYCLIST")
6 # Filter out stationary bicyclists to get only
7 # moving ones
8 moving_bicyclists = scenario_not(stationary)(
9     bicyclists, log_dir)
10 # Find bicyclists that have at least 2 other
11 # bicyclists within 5 meters (which makes a
12 # group of at least 3 bicyclists)
13 bicyclists_in_groups = near_objects(
14     moving_bicyclists, moving_bicyclists, log_dir,
15     distance_thresh=5, min_objects=2,
16     include_self=False)
17 # Output the scenario
18 output_scenario(bicyclists_in_groups, "group of
19 # at least 3 moving bicyclists within 5 meters
20 # of each other", log_dir, output_dir)

```

*Prompt: Accelerating vehicle changing lanes to the right*

```

1 # Description: accelerating vehicle changing
2 # lanes to the right
3 # First, get all vehicles in the scene
4 vehicles = get_objects_of_category(log_dir,
5     category="VEHICLE")
6 # Find vehicles that are accelerating using the
7 # accelerating function with a reliable minimum
8 # acceleration threshold of 0.65
9 accelerating_vehicles = accelerating(vehicles,
10     log_dir, min_accel=0.65)
11 # Find vehicles that are changing lanes to the
12 # right. We use the changing_lanes function
13 # with direction='right'
14 right_lane_changes = changing_lanes(vehicles,
15     log_dir, direction='right')
16 # Now find the intersection of accelerating
17 # vehicles and those changing lanes to the
18 # right. This gives us vehicles that are both
19 # accelerating and changing lanes to the right
20 accelerating_and_changing_right = scenario_and([
21     accelerating_vehicles, right_lane_changes])
22 # Output the scenario
23 output_scenario(accelerating_and_changing_right,
24     "accelerating vehicle changing lanes to the right",
25     log_dir, output_dir)

```

We provide examples of RefProg's incorrectly synthesized programs below.




Figure 8. **RefAV Object Distribution.** RefAV includes referred objects in all directions up to 150m away from the ego vehicle. RefAV places a special focus on the ego vehicle and objects that interact with the ego vehicle. Therefore, referred objects are disproportionately located on the road in front of the ego vehicle. Both referred object heatmaps are in the ego vehicle coordinate frame.

*Prompt: bus in front of ego vehicle obstructing view of red traffic light*

The following program fails because the pre-defined atomic functions lack the expressivity to synthesize a correct program. Interestingly, the LLM realizes that the provided functions don't allow us to determine if a traffic light is being obstructed.

```

1 # Description: bus in front of ego vehicle
  obstructing view of red traffic light
2
3 # Get the ego vehicle and buses
4 ego_vehicle = get_objects_of_category(log_dir,
  category="EGO_VEHICLE")
5 buses = get_objects_of_category(log_dir, category
  ="BUS")
6
7 # Find buses that are in front of the ego vehicle
8 # We need to find buses that are directly in
  front of the ego vehicle within a reasonable
  distance (20 meters) and with minimal lateral
  offset
9 buses_in_front_of_ego =
  get_objects_in_relative_direction(ego_vehicle
  , buses, log_dir, direction='forward',
  min_number=1, within_distance=20,
  lateral_thresh=3)
10
11 # We want buses that are moving in the same
  direction as the ego vehicle to ensure they'
  re actually in the same lane and potentially
  blocking the view
12 buses_in_same_direction =
  heading_in_relative_direction_to(
  buses_in_front_of_ego, ego_vehicle, log_dir,
  direction='same')
13
14 # We need buses that are in the same lane as the
  ego vehicle

```

```

15 buses_in_same_lane = in_same_lane(
  buses_in_same_direction, ego_vehicle, log_dir
  )
16
17 # These buses are likely to be obstructing the
  view of traffic lights
18 # Note: The functions don't directly allow us to
  determine if a traffic light is being
  obstructed, but buses directly in front of
  the ego vehicle in the same lane are the most
  likely candidates for obstructing the view
  of traffic lights
19
20 output_scenario(buses_in_same_lane, "bus in front
  of ego vehicle obstructing view of red
  traffic light", log_dir, output_dir)

```

*Prompt: bicycle following closely behind a vehicle*

The following program fails due to imprecise semantics because the bicycle is not necessarily following the vehicle.

```

1 # Description: bicycle following closely behind a
  vehicle
2 # We need to find bicycles that are following
  vehicles closely
3
4 # Get all bicycles and vehicles from the dataset
5 bicycles = get_objects_of_category(log_dir,
  category="BICYCLE")
6 vehicles = get_objects_of_category(log_dir,
  category="VEHICLE")
7
8 # We need to identify bicycles that are behind
  vehicles using
  get_objects_in_relative_direction to find
  bicycles that are behind vehicles. The
  bicycles should be within a reasonable
  distance (e.g., 10 meters) to be considered "
  following closely"
9 bicycles_behind_vehicles =

```

```

    get_objects_in_relative_direction(vehicles,
    bicycles, log_dir, direction="backward",
    within_distance=10, lateral_thresh=2)
10
11 # We might also want to ensure the bicycle is
    moving (not stationary)
12 moving_bicycles = scenario_not(stationary)(
    bicycles, log_dir)
13
14 # Combine the conditions: bicycles that are both
    moving and following closely behind vehicles
15 bicycles_following_vehicles = scenario_and([
    bicycles_behind_vehicles, moving_bicycles])
16
17 # Output the scenario
18 output_scenario(bicycles_following_vehicles, "
    bicycle following closely behind a vehicle",
    log_dir, output_dir)

```

*Prompt: vehicle starting to switch lanes but then going back*

The following program uses incorrect logic. It fails to combine lane changes and acceleration and uses AND on exclusive sets.

```

1 # Strategy:
2 # 1. Get all vehicles
3 # 2. Find vehicles that are changing lanes
4 # 3. Find vehicles that are also changing lanes
    in the opposite direction shortly after
5 # 4. This indicates a vehicle that started to
    change lanes but then went back
6
7 # Get all vehicles
8 vehicles = get_objects_of_category(log_dir,
    category="VEHICLE")
9
10 # Find vehicles that are moving (not stationary)
11 moving_vehicles = scenario_not(stationary)(
    vehicles, log_dir)
12
13 # Find vehicles that are changing lanes in either
    direction
14 lane_changing_vehicles = changing_lanes(
    moving_vehicles, log_dir)
15
16 # Find vehicles that are changing lanes to the
    left
17 left_lane_changes = changing_lanes(
    moving_vehicles, log_dir, direction='left')
18
19 # Find vehicles that are changing lanes to the
    right
20 right_lane_changes = changing_lanes(
    moving_vehicles, log_dir, direction='right')
21
22 # A vehicle that starts to switch lanes but goes
    back would show lateral acceleration in one
    direction followed by lateral acceleration in
    the opposite direction. We can look for
    vehicles with significant lateral
    acceleration in both directions
23
24 # Find vehicles with lateral acceleration to the
    left
25 left_accel = has_lateral_acceleration(

```

```

    moving_vehicles, log_dir, min_accel=0.5)
26
27 # Find vehicles with lateral acceleration to the
    right
28 right_accel = has_lateral_acceleration(
    moving_vehicles, log_dir, min_accel=-0.5,
    max_accel=0)
29
30 # Vehicles that have both left and right
    acceleration patterns are likely the ones
    that started to change lanes but went back
31 lane_change_aborted = scenario_and([left_accel,
    right_accel])
32 # Output the scenario
33 output_scenario(lane_change_aborted, description,
    log_dir, output_dir)

```

## J. API Listing

We present our full Python API listing below.

```

1 def has_objects_in_relative_direction(
2     track_candidates:dict,
3     related_candidates:dict,
4     log_dir:Path,
5     direction:Literal["forward", "backward", "
    left", "right"],
6     min_number:int=1,
7     max_number:int=np.inf,
8     within_distance:float=50,
9     lateral_thresh:float=np.inf)-> dict:
10     """
11     Identifies tracked objects with at least the
    minimum number of related candidates in the
    specified direction. If the minimum number is
    met, will create relationships equal to the
    max_number of closest objects.
12
13     Args:
14         track_candidates: Tracks to analyze (
    scenario dictionary). related_candidates:
    Candidates to check for in direction (
    scenario dictionary).
15         log_dir: Path to scenario logs.
16         direction: Direction to analyze from the
    track's point of view ('forward', 'backward',
    'left', 'right').
17         min_number: Minimum number of objects to
    identify in the direction per timestamp.
    Defaults to 1.
18         max_number: Maximum number of objects to
    identify in the direction per timestamp.
    Defaults to infinity.
19         within_distance: Maximum distance for
    considering an object in the direction.
    Defaults to infinity.
20         lateral_thresh: Maximum lateral distance
    the related object can be from the sides of
    the tracked object. Defaults to infinity.
21
22     Returns:
23         dict: A scenario dictionary where keys
    are track UUIDs and values are dictionaries
    containing related candidate UUIDs and lists
    of timestamps when the condition is met for
    that relative direction.

```

```

24
25 Example:
26     vehicles_with_peds_in_front =
27     has_objects_in_relative_direction(vehicles,
28     pedestrians, log_dir, direction='forward',
29     min_number=2)
30     """
31 def get_objects_in_relative_direction(
32     track_candidates:dict,
33     related_candidates:dict,
34     log_dir:Path,
35     direction:Literal["forward", "backward", "
36     left", "right"],
37     min_number:int=0,
38     max_number:int=np.inf,
39     within_distance:float=50,
40     lateral_thresh:float=np.inf)->dict:
41     """
42     Returns a scenario dictionary of the related
43     candidates that are in the relative direction
44     of the track candidates.
45
46     Args:
47     track_candidates: Tracks (scenario
48     dictionary).
49     related_candidates: Candidates to check
50     for in direction (scenario dictionary).
51     log_dir: Path to scenario logs.
52     direction: Direction to analyze from the
53     track's point of view ('forward', 'backward',
54     'left', 'right').
55     min_number: Minimum number of objects to
56     identify in the direction per timestamp.
57     Defaults to 0.
58     max_number: Maximum number of objects to
59     identify in the direction per timestamp.
60     Defaults to infinity.
61     within_distance: Maximum distance for
62     considering an object in the direction.
63     Defaults to infinity.
64     lateral_thresh: Maximum lateral distance
65     the related object can be from the sides of
66     the tracked object. Lateral distance is
67     distance is the distance from the sides of
68     the object that are parallel to the specified
69     direction. Defaults to infinity.
70
71     Returns:
72     dict: A scenario dictionary where keys
73     are track UUIDs and values are dictionaries
74     containing related candidate UUIDs and lists
75     of timestamps when the condition is met for
76     that relative direction.
77
78     Example:
79     peds_in_front_of_vehicles =
80     get_objects_in_relative_direction(vehicles,
81     pedestrians, log_dir, direction='forward',
82     min_number=2)
83     """
84
85 def get_objects_of_category(
86     log_dir,

```

```

63     category)->dict:
64     """
65     Returns all objects from a given category
66     from the log annotations. This method accepts
67     the super-categories "ANY" and "VEHICLE".
68
69     Args:
70     log_dir: Path to the directory containing
71     scenario logs and data.
72     category: the category of objects to
73     return
74
75     Returns:
76     dict: A scenario dict that where keys are
77     the unique id (uuid) of the object and
78     values are the list of timestamps the object
79     is in view of the ego-vehicle.
80
81     Example:
82     trucks = get_objects_of_category(log_dir,
83     category='TRUCK')
84     """
85
86 def is_category(
87     track_candidates:dict,
88     log_dir:Path,
89     category:str)->dict:
90     """
91     Returns all objects from a given category
92     from track_candidates dict. This method
93     accepts the super-categories "ANY" and "
94     VEHICLE".
95
96     Args:
97     track_candidates: The scenario dict
98     containing the objects to filter down
99     log_dir: Path to the directory containing
100     scenario logs and data.
101     category: the category of objects to
102     return
103
104     Returns:
105     dict: A scenario dict that where keys are
106     the unique id of the object of the given
107     category and values are the list of
108     timestamps the object is in view of the ego-
109     vehicle.
110
111     Example:
112     box_trucks = is_category(vehicles,
113     log_dir, category='BOX_TRUCK')
114     """
115
116 def turning(
117     track_candidates: dict,
118     log_dir:Path,
119     direction:Literal["left", "right", None]=None
120 )->dict:
121     """
122     Returns objects that are turning in the given
123     direction.
124
125     Args:
126     track_candidates: The objects you want to
127     filter from (scenario dictionary).

```

```

108     log_dir: Path to scenario logs.
109     direction: The direction of the turn,
    from the track's point of view ('left', '
    right', None).
110
111 Returns:
112     dict: A filtered scenario dictionary
    where keys are track UUIDs that meet the
    turning criteria and values are nested
    dictionaries containing timestamps.
113
114 Example:
115     turning_left = turning(vehicles, log_dir,
    direction='left')
116     """
117
118 def changing_lanes(
119     track_candidates:dict,
120     log_dir:Path,
121     direction:Literal["left", "right", None]=None
122 )-> dict:
123     """
124     Identifies lane change events for tracked
    objects in a scenario.
125
126 Args:
127     track_candidates: The tracks to analyze (
    scenario dictionary).
128     log_dir: Path to scenario logs.
129     direction: The direction of the lane
    change. None indicates tracking either left
    or right lane changes ('left', 'right', None)
    .
130
131 Returns:
132     dict: A filtered scenario dictionary
    where keys are track UUIDs that meet the lane
    change criteria and values are nested
    dictionaries containing timestamps and
    related data.
133
134 Example:
135     left_lane_changes = changing_lanes(
    vehicles, log_dir, direction='left')
136     """
137
138
139 def has_lateral_acceleration(
140     track_candidates:dict,
141     log_dir:Path,
142     min_accel=-np.inf,
143     max_accel=np.inf)-> dict:
144     """
145     Objects with a lateral acceleration between
    the minimum and maximum thresholds. Most
    objects with a high lateral acceleration are
    turning. Postive values indicate acceleration
    to the left while negative values indicate
    acceleration to the right.
146
147 Args:
148     track_candidates: The tracks to analyze (
    scenario dictionary).
149     log_dir: Path to scenario logs.
150     direction: The direction of the lane
    change. None indicates tracking either left

```

```

    or right lane changes ('left', 'right', None)
    .
151
152 Returns:
153     dict: A filtered scenario dictionary
    where keys are track UUIDs that meet the lane
    change criteria and values are nested
    dictionaries containing timestamps and
    related data.
154
155 Example:
156     jerking_left = has_lateral_acceleration(
    non_turning_vehicles, log_dir, min_accel=2)
157     """
158
159
160 def facing_toward(
161     track_candidates:dict,
162     related_candidates:dict,
163     log_dir:Path,
164     within_angle:float=22.5,
165     max_distance:float=50)->dict:
166     """
167     Identifies objects in track_candidates that
    are facing toward objects in related
    candidates. The related candidate must lie
    within a region lying within within_angle
    degrees on either side the track-candidate's
    forward axis.
168
169 Args:
170     track_candidates: The tracks that could
    be heading toward another tracks
171     related_candidates: The objects to
    analyze to see if the track_candidates are
    heading toward
172     log_dir: Path to the directory
    containing scenario logs and data.
173     fov: The field of view of the
    track_candidates. The related candidate must
    lie within a region lying within fov/2
    degrees on either side the track-candidate's
    forward axis.
174     max_distance: The maximum distance a
    related_candidate can be away to be
    considered by
175
176 Returns:
177     dict: A filtered scenario dict that
    contains the subset of track candidates
    heading toward at least one of the related
    candidates.
178
179 Example:
180     pedestrian_facing_away = scenario_not(
    facing_toward)(pedestrian, ego_vehicle,
    log_dir, within_angle=180)
181     """
182
183
184 def heading_toward(
185     track_candidates:dict,
186     related_candidates:dict,
187     log_dir:Path,
188     angle_threshold:float=22.5,
189     minimum_speed:float=.5,
190     max_distance:float=np.inf)->dict:

```

```

191 """
192 Identifies objects in track_candidates that
    are heading toward objects in related
    candidates. The track candidates acceleration
    vector must be within the given angle
    threshold of the relative position vector.
    The track candidates must have a component of
    velocity toward the related candidate
    greater than the minimum_accel.
193
194 Args:
195     track_candidates: The tracks that could
    be heading toward another tracks
196     related_candidates: The objects to
    analyze to see if the track_candidates are
    heading toward
197     log_dir: Path to the directory
    containing scenario logs and data.
198     angle_threshold: The maximum angular
    difference between the velocity vector and
    relative position vector between the track
    candidate and related candidate.
199     min_vel: The minimum magnitude of the
    component of velocity toward the related
    candidate
200     max_distance: Distance in meters the
    related candidates can be away from the track
    candidate to be considered
201
202 Returns:
203     dict: A filtered scenario dict that
    contains the subset of track candidates
    heading toward at least one of the related
    candidates.
204
205 Example:
206     heading_toward_traffic_cone =
    heading_toward(vehicles, traffic_cone,
    log_dir)
207 """
208
209 def accelerating(
210     track_candidates:dict,
211     log_dir:Path,
212     min_accel:float=.65,
213     max_accel:float=np.inf)->dict:
214     """
215     Identifies objects in track_candidates that
    have a forward acceleration above a threshold
    . Values under -1 reliably indicates braking.
    Values over 1.0 reliably indicates
    accelerating.
216
217
218 Args:
219     track_candidates: The tracks to analyze
    for acceleration (scenario dictionary)
220     log_dir: Path to the directory
    containing scenario logs and data.
221     min_accel: The lower bound of
    acceleration considered
222     max_accel: The upper bound of
    acceleration considered
223
224 Returns:
225     dict: A filtered scenario dictionary
    containing the objects with an acceleration

```

```

    between the lower and upper bounds.
226
227 Example:
228     accelerating_motorcycles = accelerating(
    motorcycles, log_dir)
229 """
230
231 def has_velocity(
232     track_candidates:dict,
233     log_dir:Path,
234     min_velocity:float=.5,
235     max_velocity:float=np.inf)->dict:
236     """
237     Identifies objects with a velocity between
    the given maximum and minimum velocities in m
    /s. Stationary objects may have a velocity up
    to 0.5 m/s due to annotation jitter.
238
239
240 Args:
241     track_candidates: Tracks to analyze (
    scenario dictionary).
242     log_dir: Path to scenario logs.
243     min_velocity: Minimum velocity (m/s).
    Defaults to 0.5.
244     max_velocity: Maximum velocity (m/s)
245
246 Returns:
247     dict: Filtered scenario dictionary of
    objects meeting the velocity criteria.
248
249 Example:
250     fast_vehicles = has_min_velocity(vehicles
    , log_dir, min_velocity=5)
251 """
252
253 def at_pedestrian_crossing(
254     track_candidates:dict,
255     log_dir:Path,
256     within_distance:float=1)->dict:
257     """
258     Identifies objects that within a certain
    distance from a pedestrian crossing. A
    distance of zero indicates that the object is
    within the boundaries of the pedestrian
    crossing.
259
260
261 Args:
262     track_candidates: Tracks to analyze (
    scenario dictionary).
263     log_dir: Path to scenario logs.
264     within_distance: Distance in meters the
    track candidate must be from the pedestrian
    crossing. A distance of zero means that the
    object must be within the boundaries of the
    pedestrian crossing.
265
266 Returns:
267     dict: Filtered scenario dictionary where
    keys are track UUIDs and values are lists of
    timestamps.
268
269 Example:
270     vehicles_at_ped_crossing =
    at_pedestrian_crossing(vehicles, log_dir)
271 """

```

```

272
273
274 def on_lane_type(
275     track_uuid:dict,
276     log_dir,
277     lane_type:Literal["BUS", "VEHICLE", "BIKE"])
278     ->dict:
279     """
280     Identifies objects on a specific lane type.
281
282     Args:
283         track_candidates: Tracks to analyze (
284             scenario dictionary).
285         log_dir: Path to scenario logs.
286         lane_type: Type of lane to check ('BUS',
287             'VEHICLE', or 'BIKE').
288
289     Returns:
290         dict: Filtered scenario dictionary where
291             keys are track UUIDs and values are lists of
292             timestamps.
293
294     Example:
295         vehicles_on_bus_lane = on_lane_type(
296             vehicles, log_dir, lane_type="BUS")
297     """
298
299 def near_intersection(
300     track_uuid:dict,
301     log_dir:Path,
302     threshold:float=5)->dict:
303     """
304     Identifies objects within a specified
305     threshold of an intersection in meters.
306
307     Args:
308         track_candidates: Tracks to analyze (
309             scenario dictionary).
310         log_dir: Path to scenario logs.
311         threshold: Distance threshold (in meters)
312             to define "near" an intersection.
313
314     Returns:
315         Filtered scenario dictionary where keys
316             are track UUIDs and values are lists of
317             timestamps.
318
319     Example:
320         bicycles_near_intersection =
321             near_intersection(bicycles, log_dir,
322                 threshold=10.0)
323     """
324
325 def on_intersection(
326     track_candidates:dict,
327     log_dir:Path)->dict:
328     """
329     Identifies objects located on top of an road
330     intersection.
331
332     Args:
333         track_candidates: Tracks to analyze (
334             scenario dictionary).
335         log_dir: Path to scenario logs.

```

```

324 Returns:
325     dict: Filtered scenario dictionary where
326         keys are track UUIDs and values are lists of
327         timestamps.
328
329 Example:
330     strollers_on_intersection =
331         on_intersection(strollers, log_dir)
332     """
333
334 def being_crossed_by(
335     track_candidates:dict,
336     related_candidates:dict,
337     log_dir:Path,
338     direction:Literal["forward", "backward", "
339         left",
340         "right"]="forward",
341     in_direction:Literal['clockwise', '
342         counterclockwise',
343         'either']='either',
344     forward_thresh:float=10,
345     lateral_thresh:float=5)->dict:
346     """
347     Identifies objects that are being crossed by
348     one of the related candidate objects. A
349     crossing is defined as the related candidate's
350     centroid crossing the half-midplane of a
351     tracked candidate. The direction of the half-
352     midplane is specified with the direction.
353
354     Args:
355         track_candidates: Tracks to analyze .
356         related_candidates: Candidates (e.g.,
357             pedestrians or vehicles) to check for
358             crossings.
359         log_dir: Path to scenario logs.
360         direction: specifies the axis and
361             direction the half midplane extends from
362         in_direction: which direction the related
363             candidate has to cross the midplane for it
364             to be considered a crossing
365         forward_thresh: how far the midplane
366             extends from the edge of the tracked object
367         lateral_thresh: the two planes offset
368             from the midplane. If an related candidate
369             crosses the midplane, it will continue being
370             considered crossing until it goes past the
371             lateral_thresh.
372
373     Returns:
374         dict: A filtered scenario dictionary
375             containing all of the track candidates that
376             were crossed by the related candidates given
377             the specified constraints.
378
379     Example:
380         overtaking_on_left = being_crossed_by(
381             moving_cars, moving_cars, log_dir, direction=
382             "left", in_direction="clockwise",
383             forward_thresh=4)
384         vehicles_crossed_by_peds =
385             being_crossed_by(vehicles, pedestrians,
386                 log_dir)
387     """

```

```

363 def near_objects (
364     track_uuid:dict,
365     candidate_uuids:dict,
366     log_dir:Path,
367     distance_thresh:float=10,
368     min_objects:int=1,
369     include_self:bool=False)->dict:
370     """
371     Identifies timestamps when a tracked object
372     is near a specified set of related objects.
373
374     Args:
375         track_candidates: Tracks to analyze (
376         scenario dictionary).
377         related_candidates: Candidates to check
378         for proximity (scenario dictionary).
379         log_dir: Path to scenario logs.
380         distance_thresh: Maximum distance in
381         meters a related candidate can be away to be
382         considered "near".
383         min_objects: Minimum number of related
384         objects required to be near the tracked
385         object.
386
387     Returns:
388         dict: A scenario dictionary where keys
389         are timestamps when the tracked object is
390         near the required number of related objects
391         and values are lists of related candidate
392         UUIDs present at those timestamps.
393
394     Example:
395         vehicles_near_ped_group = near_objects(
396         vehicles, pedestrians, log_dir, min_objects
397         =3)
398     """
399
400 def following(
401     track_uuid:dict,
402     candidate_uuids:dict,
403     log_dir:Path)-> dict:
404     """
405     Returns timestamps when the tracked object is
406     following a lead object. Following is
407     defined simultaneously moving in the same
408     direction and lane.
409     """
410
411 def heading_in_relative_direction_to(
412     track_candidates,
413     related_candidates,
414     log_dir,
415     direction:Literal['same', 'opposite', '
416     perpendicular'])->dict:
417     """
418     Returns the subset of track candidates that
419     are traveling in the given direction compared
420     to the related candidates.
421
422     Arguments:
423         track_candidates: The set of objects that
424         could be traveling in the given direction
425         related_candidates: The set of objects
426         that the direction is relative to
427         log_dir: The path to the log data

```

```

428     direction: The direction that the
429     positive tracks are
430     traveling in relative to the related
431     candidates:
432         - "opposite" indicates the track
433         candidates are traveling in a direction
434         135-180 degrees from the direction the
435         related candidates are heading toward.
436         - "same" indicates the track
437         candidates that are traveling in a direction
438         0-45 degrees from the direction the related
439         candidates are heading toward.
440         - "perpendicular" indicates the track
441         candidates that are traveling in a direction
442         45-135 degrees from the direction the
443         related candiates are heading toward.
444
445     Returns:
446         dict: the subset of track candidates that
447         are traveling in the given direction
448         compared to the related candidates.
449
450     Example:
451         oncoming_traffic =
452         heading_in_relative_direction_to(vehicles,
453         ego_vehicle, log_dir, direction='opposite')
454     """
455
456 def stationary(
457     track_candidates:dict,
458     log_dir:Path)->dict:
459     """
460     Returns objects that moved less than 2m over
461     their length of observation in the scneario.
462     This object is only intended to separate
463     parked from active vehicles. Use has_velocity
464     () with thresholding if you want to indicate
465     vehicles that are temporarily stopped.
466
467     Args:
468         track_candidates: Tracks to analyze (
469         scenario dictionary).
470         log_dir: Path to scenario logs.
471
472     Returns:
473         dict: A filtered scenario dictionary
474         where keys are track UUIDs and values are
475         lists of timestamps when the object is
476         stationary.
477
478     Example:
479         parked_vehicles = stationary(vehicles,
480         log_dir)
481     """
482
483 def at_stop_sign(
484     track_candidates:dict,
485     log_dir:Path,
486     forward_thresh:float=10)->dict:
487     """
488     Identifies timestamps when a tracked object
489     is in a lane corresponding to a stop sign.
490     The tracked object must be within 15m of the
491     stop sign. This may highlight vehicles using
492     street parking near a stopped sign.

```

```

447
448     Args:
449         track_candidates: Tracks to analyze (
450         scenario dictionary).
451         log_dir: Path to scenario logs.
452         forward_thresh: Distance in meters the
453         vehicle is from the stop sign in the stop
454         sign's front direction
455
456     Returns:
457         dict: A filtered scenario dictionary
458         where keys are track UUIDs and values are
459         lists of timestamps when the object is at a
460         stop sign.
461
462     Example:
463         vehicles_at_stop_sign = at_stop_sign(
464         vehicles, log_dir)
465     """
466
467 def in_drivable_area(
468     track_candidates:dict,
469     log_dir:Path)->dict:
470     """
471     Identifies objects within track_candidates
472     that are within a drivable area.
473
474     Args:
475         track_candidates: Tracks to analyze (
476         scenario dictionary).
477         log_dir: Path to scenario logs.
478
479     Returns:
480         dict: A filtered scenario dictionary
481         where keys are track UUIDs and values are
482         lists of timestamps when the object is in a
483         drivable area.
484
485     Example:
486         buses_in_drivable_area = in_drivable_area
487         (buses, log_dir)
488     """
489
490 def on_road(
491     track_candidates:dict,
492     log_dir:Path)->dict:
493     """
494     Identifies objects that are on a road or bike
495     lane. This function should be used in place
496     of in_drivable_area() when referencing
497     objects that are on a road. The road does not
498     include parking lots or other driveable
499     areas connecting the road to parking lots.
500
501     Args:
502         track_candidates: Tracks to filter (
503         scenario dictionary).
504         log_dir: Path to scenario logs.
505
506     Returns:
507         dict: The subset of the track candidates
508         that are currently on a road.
509
510     Example:
511         animals_on_road = on_road(animals,

```

```

512     log_dir)
513     """
514
515 def in_same_lane(
516     track_candidates:dict,
517     related_candidates:dict,
518     log_dir:Path)->dict:
519     """
520     Identifies tracks that are in the same road
521     lane as a related candidate.
522
523     Args:
524         track_candidates: Tracks to filter (
525         scenario dictionary)
526         related_candidates: Potential objects
527         that could be in the same lane as the track (
528         scenario dictionary)
529         log_dir: Path to scenario logs.
530
531     Returns:
532         dict: A filtered scenario dictionary
533         where keys are track UUIDs and values are
534         lists of timestamps when the object is on a
535         road lane.
536
537     Example:
538         bicycle_in_same_lane_as_vehicle =
539         in_same_lane(bicycle, regular_vehicle,
540         log_dir)
541     """
542
543 def on_relative_side_of_road(
544     track_candidates:dict,
545     related_candidates:dict,
546     log_dir:Path,
547     side=Literal['same', 'opposite'])->dict:
548     """
549     Identifies tracks that are in the same road
550     lane as a related candidate.
551
552     Args:
553         track_candidates: Tracks to filter (
554         scenario dictionary)
555         related_candidates: Potential objects
556         that could be in the same lane as the track (
557         scenario dictionary)
558         log_dir: Path to scenario logs.
559
560     Returns:
561         dict: A filtered scenario dictionary
562         where keys are track UUIDs and values are
563         lists of timestamps when the object is on a
564         road lane.
565
566     Example:
567         bicycle_in_same_lane_as_vehicle =
568         in_same_lane(bicycle, regular_vehicle,
569         log_dir)
570     """
571
572 def is_color(
573     track_candidates: dict,
574     log_dir: Path,
575     color:Literal["white", "silver", "black", "
576     red", "yellow", "blue"])->dict:
577     """

```

```

541 Returns objects that are the given color,
542 determined by SIGLIP2 feature similarity.
543
544 Args:
545     track_candidates: The objects you want to
546     filter from (scenario dictionary).
547     log_dir: Path to scenario logs.
548     color: The color of the objects you want
549     to return. Must be one of 'white', 'silver',
550     'black', 'red', 'yellow', or 'blue'.
551     Inputting a different color defaults to
552     returning all objects.
553
554 Returns:
555     dict: A filtered scenario dictionary
556     where keys are track UUIDs that meet the
557     turning criteria and values are nested
558     dictionaries containing timestamps.
559
560 Example:
561     red_cars = is_color(cars, log_dir, color
562     ='red')
563     """
564
565 def scenario_and(scenario_dicts)->dict:
566     """
567     Returns a composed scenario where the track
568     objects are the intersection of all of the
569     track objects with the same uuid and
570     timestamps.
571
572     Args:
573         scenario_dicts: the scenarios to combine
574
575     Returns:
576         dict: a filtered scenario dictionary that
577         contains tracked objects found in all given
578         scenario dictionaries
579
580     Example:
581         jaywalking_peds = scenario_and([
582         peds_on_road, peds_not_on_pedestrian_crossing
583         ])
584         """
585
586 def scenario_or(scenario_dicts)->dict:
587     """
588     Returns a composed scenario where that tracks
589     all objects and relationships in all of the
590     input scenario dicts.
591
592     Args:
593         scenario_dicts: the scenarios to combine
594
595     Returns:
596         dict: an expanded scenario dictionary
597         that contains every tracked object in the
598         given scenario dictionaries
599
600     Example:
601         be_cautious_around = scenario_or([
602         animal_on_road, stroller_on_road])
603         """
604
605 def reverse_relationship(func)->dict:

```

```

586     """
587     Wraps relational functions to switch the top
588     level tracked objects and relationships
589     formed by the function.
590
591     Args:
592         relational_func: Any function that takes
593         track_candidates and related_candidates as
594         its first and second arguments
595
596     Returns:
597         dict: scenario dict with swapped top-
598         level tracks and related candidates
599
600     Example:
601         group_of_peds_near_vehicle =
602         reverse_relationship(near_objects)(vehicles,
603         peds, log_dir, min_objects=3)
604         """
605
606 def scenario_not(func)->dict:
607     """
608     Wraps composable functions to return the
609     difference of the input track dict and output
610     scenario dict. Using scenario_not with a
611     composable relational function will not
612     return any relationships.
613
614     Args:
615         composable_func: Any function that takes
616         track_candidates as its first input
617
618     Returns:
619         dict: scenario dict with tracks that are
620         not filtered by func
621
622     Example:
623         active_vehicles = scenario_not(stationary
624         )(vehicles, log_dir)
625         """
626
627 def output_scenario(
628     scenario:dict,
629     description:str,
630     log_dir:Path,
631     output_dir:Path,
632     visualize:bool=False,
633     **visualization_kwargs):
634     """
635     Outputs a file containing the predictions in
636     an evaluation-ready format. Do not provide
637     any visualization kwargs.
638     """

```