

KnowVal: A Knowledge-Augmented and Value-Guided Autonomous Driving System

Supplementary Material

A. Implementation Details

A.1. Training and Architecture Details

We employ Qwen2.5-3B for retrieval and knowledge embedding, with the number of retrieved knowledge entries set to $N_K = 16$. For the planner, the number of candidate trajectories is $N_T = 20$, and the Value Model performs $L = 3$ iterative reasoning steps. The decay factor used in the total value assessment is $\gamma = 0.7$. To increase trajectory diversity, baseline models on nuScenes are fine-tuned for three epochs using eight A100 GPUs, while models on Bench2Drive are fine-tuned for one epoch under the same configuration. The Value Model is trained separately on eight V100 GPUs for 50 epochs using the AdamW optimizer with a cosine annealing learning rate schedule.

A.2. Prompt used for scoring trajectories with MLLM

```
# ROLE
You are an expert autonomous driving
auditor. Your analysis is meticulous,
objective, and follows a strict logical
process. Your primary goal is to
identify and quantify risk.

# TASK
For each of the prioritized driving rules
provided below, you MUST follow a
mandatory **4-Step Critical Evaluation
Process**. Your final score for each
rule is a direct consequence of this
process. Do not skip any steps.

# PRIORITIZED RULES TO EVALUATE
{formatted_rules}

# MANDATORY 4-STEP CRITICAL EVALUATION
PROCESS
For each rule, you must internally answer
these questions in order:

**Step 1: Applicability Check**
* **Question:** Is this rule relevant to
the current scene? Can it possibly be
applied or violated based on the visual
context?
* **Outcome:**
  * If **NO** (e.g., a traffic light
rule when no traffic lights are visible)
, the process stops here. Your
```

```
conclusion is "Not Applicable". **The
score MUST be 1.0.**
  * If **YES**, proceed to Step 2.

**Step 2: Evidence Check**
* **Question:** Does the image provide
any visual information (positive or
negative) to judge adherence to this
applicable rule?
* **Outcome:**
  * If **NO** (e.g., a rule about
checking blind spots, but the driver's
actions are not visible), the process
stops here. Your conclusion is "
Applicable, but no evidence". **The
score MUST be 0.0.**
  * If **YES**, proceed to Step 3.

**Step 3: Adherence Check**
* **Question:** Based on the available
evidence, is the ego vehicle fully and
correctly adhering to the rule?
* **Outcome:**
  * If **YES**, the process stops here.
Your conclusion is "Perfect Adherence".
**The score MUST be 1.0.**
  * If **NO**, there is a violation.
Proceed to Step 4 to determine its
severity.

**Step 4: Risk Assessment (FOR VIOLATIONS
ONLY)**
* **Question:** How severe is the
immediate risk created by this violation
?
* **Outcome:** Assign a risk level, which
then determines the score range.
  * **Negligible Risk:** A technical
error with no immediate safety impact (e
.g., stopping 10cm past the stop line
with no cross-traffic). **Score: -0.1 to
-0.2**.
  * **Low Risk:** A poor driving habit
or minor violation that increases
potential risk but doesn't create
immediate danger (e.g., late signaling
in light traffic). **Score: -0.3 to
-0.4**.
  * **Moderate Risk:** A clear unsafe
action that significantly increases risk
(e.g., tailgating, changing lanes
```

without signal near other cars). **Score : -0.5 to -0.7**.

- * **High Risk:** A critical error that creates a probable chance of collision or puts others in immediate danger (e.g., failing to yield to a pedestrian entering a crosswalk, running a stale yellow light). **Score: -0.8 to -1.0**.

CRITICAL INSTRUCTIONS

- **AVOID LAZY SCORING:** Do not default to 1.0 or 0.0 unless the 4-step process explicitly leads you there. The goal is to differentiate between different levels of non-compliance.
- **JUSTIFY YOUR SCORE:** The 'conclusion' and the 'risk_level' must logically and transparently lead to the final 'score'.
- **BE SPECIFIC:** In your evidence description, refer to specific objects and actions in the scene.

OUTPUT FORMAT

You MUST respond with a single JSON object. The object must have a single key "evaluations", which is a list. Each item in the list must adhere to the following structure:

1. A "rule" key with the original rule text.
2. A "reasoning" object containing four keys:
 - "positive_evidence": Visual cues that the rule is being followed. State "None" if none.
 - "negative_evidence": Visual cues that the rule is being violated. State "None" if none.
 - "risk_level": A string describing the risk for violations ("Negligible", "Low", "Moderate", "High"). For non-violations, state "N/A".
 - "conclusion": A summary of your 4-step analysis, explicitly stating the final judgment (e.g., "Not Applicable", "Perfect Adherence", "Moderate Risk Violation").
3. A "score" key, a float number that is a direct consequence of your 'conclusion' and 'risk_level'.

Example of your required output format:

```
```json
{
 "scene_description": "A vehicle is approaching an intersection with a pedestrian waiting to cross at a marked
```

```
crosswalk. The traffic light is green.",
"evaluations": [
 {
 "rule": "Yield to pedestrians at the crosswalk.",
 "reasoning": {
 "positive_evidence": "The vehicle is slowing down as it approaches the crosswalk.",
 "negative_evidence": "The vehicle does not come to a complete stop and proceeds through the crosswalk, forcing the pedestrian to wait.",
 "risk_level": "High",
 "conclusion": "High Risk Violation. The rule is applicable and evidence is clear. The vehicle failed to yield to a pedestrian ready to cross, creating immediate danger."
 },
 "score": -0.8
 },
 {
 "rule": "Obey traffic signals.",
 "reasoning": {
 "positive_evidence": "The traffic light for the vehicle's direction of travel is green, and the vehicle is proceeding through the intersection.",
 "negative_evidence": "None",
 "risk_level": "N/A",
 "conclusion": "Perfect Adherence. The rule is applicable, and the vehicle is correctly following the green light signal."
 },
 "score": 1.0
 },
 {
 "rule": "No Honking.",
 "reasoning": {
 "positive_evidence": "None",
 "negative_evidence": "None",
 "risk_level": "N/A",
 "conclusion": "Applicable, but no evidence. The rule is relevant in a city scene, but there is no visual or contextual information to determine if the horn was used."
 },
 "score": 0.0
 },
 {
 "rule": "Stop at Railroad Crossing.",
 "reasoning": {
 "positive_evidence": "None",
 "negative_evidence": "None",
 "risk_level": "N/A",
```

```

 "conclusion": "Not Applicable.
 There are no railway tracks or crossings
 visible in the scene."
 }},
 "score": 1.0
 }}
}
}}""

```

Listing 1. The prompt used for constructing GT of the Value Model Dataset

### A.3. Prompt for keyword extraction in the retrieval stage

```

---Goal---
Given a text document that is potentially
relevant to this activity and a list of
entity types, identify all entities of
those types from the text and all
relationships among the identified
entities.
Use {language} as output language.

---Steps---
1. Identify all entities. For each
 identified entity, extract the following
 information:
- entity_name: Name of the entity, use same
 language as input text. If English,
 capitalized the name.
- entity_type: One of the following types:
 [{entity_types}]
- entity_description: Comprehensive
 description of the entity's attributes
 and activities
Format each entity as ("entity"{
 tuple_delimiter}<entity_name>{
 tuple_delimiter}<entity_type>{
 tuple_delimiter}<entity_description>)

2. From the entities identified in step 1,
 identify all pairs of (source_entity,
 target_entity) that are *clearly related
 * to each other.
For each pair of related entities, extract
the following information:
- source_entity: name of the source entity,
 as identified in step 1
- target_entity: name of the target entity,
 as identified in step 1
- relationship_description: explanation as
 to why you think the source entity and
 the target entity are related to each
 other
- relationship_strength: a numeric score
 indicating strength of the relationship

```

```

between the source entity and target
entity
- relationship_keywords: one or more high-
level key words that summarize the
overarching nature of the relationship,
focusing on concepts or themes rather
than specific details
Format each relationship as ("relationship
"{tuple_delimiter}<source_entity>{
tuple_delimiter}<target_entity>{
tuple_delimiter}<
relationship_description>{
tuple_delimiter}<relationship_keywords>{
tuple_delimiter}<relationship_strength>)

3. Identify high-level key words that
summarize the main concepts, themes, or
topics of the entire text. These should
capture the overarching ideas present in
the document.
Format the content-level key words as ("
content_keywords"{tuple_delimiter}<
high_level_keywords>)

4. Return output in {language} as a single
list of all the entities and
relationships identified in steps 1 and
2. Use **{record_delimiter}** as the
list delimiter.

5. When finished, output {
completion_delimiter}

#####
---Examples---
#####
{examples}

#####
---Real Data---
#####
Entity_types: [{entity_types}]
Text:
{input_text}
#####
Output:

```

Listing 2. The prompt used for retrieval

## B. About the Knowledge Graph

A local view of the knowledge graph is shown in Figure 7. Currently, the knowledge graph contains 1,324 nodes and 2,785 edges.

## C. Additional Ablation

**Ablation on nuScenes.** As shown in Table 6, integrating the knowledge retrieval and value model components, as well as incorporating open-world perception and retrieval-guided supplementary perception, consistently reduces the collision rate. Although the deviation from human driving trajectories increases slightly, this does not imply a decline in planning quality. We further analyze the effects of varying the numbers of knowledge entries and candidate trajectories. As shown in the middle and lower sections of the table, increasing these quantities beyond a certain threshold provides minimal additional benefit, indicating diminishing returns.

**Ablation on Bench2Drive.** As shown in Table 7, The experiments show each component’s effectiveness and robustness to noisy knowledge items.

**Ablation on the Value Model.** As shown in Table 8, we evaluate several architectures for the Value Model. Using BEV images as input processed by CNNs offers the advantage of not requiring model-specific training for each baseline, but yields weaker predictive performance. Introducing weighted averaging across disaggregated subscores effectively reduces prediction errors. The entry labeled *10K*

represents our initial dataset version, which included only trajectories generated by HENet++. In the final version, we augmented the dataset with numerous randomly generated low-quality trajectories as negative samples, significantly increasing both scale and diversity.

## D. Analysis of Retrieval Validity

We conducted statistics on nuScenes *val*. **50.1%** of the retrieved items are Key Items (which certainly fit in this scenario), **41.8%** of them are General Items (fit in any scenario, e.g., speed limits), and **8.1%** of them are Irrelevant Items (e.g., ‘give way to ambulance’, but there is no ambulance). Crucially, the fraction of Irrelevant Items did not exceed 2/16 in any single trajectory. **The rows H to J in Table 7 also demonstrate robustness to noisy entries**, as we train the value model to score 0 for irrelevant entries.

## E. Qualitative Analysis of the Value Model

To better illustrate the role of our proposed and trained Value Model, Figure 8 showcases several examples and visualizes ego vehicle trajectories across two scenarios, where trajectories 2 and 4 were manually constructed as *negative samples* and the black trajectory in the center repre-

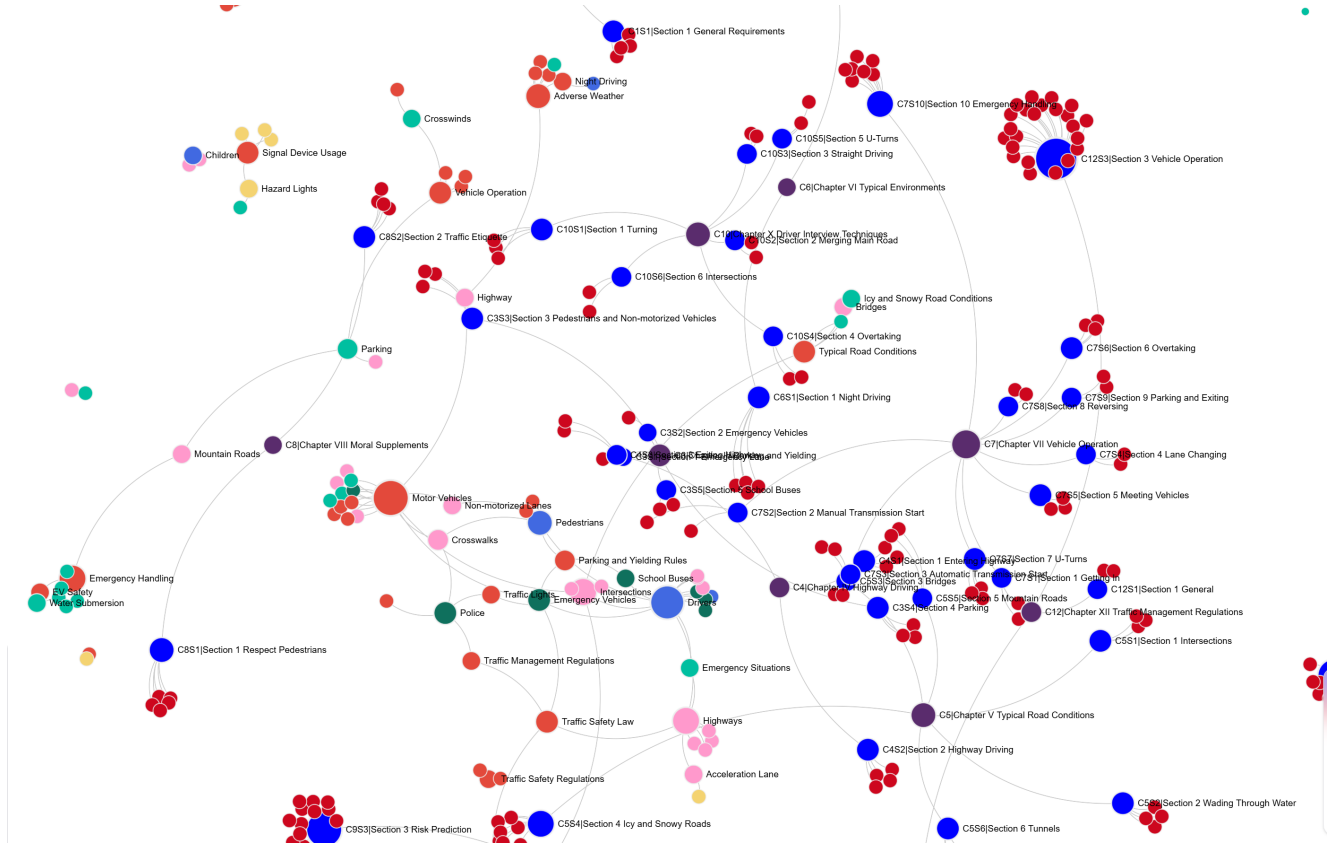


Figure 7. Local view of the knowledge graph.

Table 6. **Ablation studies on nuScenes.** The baseline model is GenAD. K&V: using the proposed knowledge retrieval and value model. OW: incorporating open-ended perception and understanding of abstract concepts. RgP: using retrieval results to guide supplementary perception.  $N_K$ : the number of retrieved knowledge entries.  $W_K$ : applying weighting to the total score.  $N_T$ : the number of candidate trajectories.

K&V	OW	RgP	$N_K$	$W_K$	$N_T$	L2(m)↓	Col↓
			16	✓	20	0.91	0.43
✓			16	✓	20	0.89	0.38
✓	✓		16	✓	20	1.47	0.34
✓	✓	✓	16	✓	20	1.51	0.33
✓	✓	✓	16	✓	10	1.56	0.37
✓	✓	✓	16	✓	20	1.51	0.33
✓	✓	✓	16	✓	30	1.52	0.33
✓	✓	✓	8	✓	20	1.29	0.36
✓	✓	✓	16		20	1.67	0.38
✓	✓	✓	16	✓	20	1.51	0.33
✓	✓	✓	24	✓	20	1.49	0.34

Table 7. **Ablation studies on Bench2Drive.** Exp. A is the SimLingo baseline. Exp. B uses the retrieved language tokens as Key-Value pairs for the planning transformer. Exp. C trains the value model using only the future-state tokens  $S_i$  as input—excluding knowledge entries—while keeping the training setting consistent. Exp. D uses heuristic rules (SparseDrive; Sun et al., 2025) rather than a learned value model. Exp. H randomly inserts  $N_k$  random knowledge items into the retrieved entries, while Exp. I uses  $N_k$  randomly selected knowledge items without any retrieval. The experiments show each component’s effectiveness and robustness to noisy knowledge items.

ID	Knowledge	Values	OW	RgP	DS	SR
A					85.07	67.27
B	✓				84.91	67.03
C		✓			85.92	67.69
D		Heuristic Rules			85.94	67.82
E	✓	✓			87.61	68.40
F	✓	✓		✓	88.16	68.85
G	✓	✓		✓	88.42	69.03
H	Noise	✓			87.07	67.99
I	Random	✓			86.18	67.74
J	✓	✓			87.61	68.40

Table 8. **Comparison on different Value Model architectures.**

Input $S_i$	Output	Traning Scale	MSE↓	MAE↓
BEV Image	Total Score	10K	0.35	0.59
BEV Image	Total Score	128K	0.22	0.44
Scene Feature	Total Score	128K	0.13	0.27
Scene Feature	Subscores	128K	0.11	0.20

tory successfully avoids pedestrians and the leading vehicle, whereas the trajectory in the second image heads directly toward them. A comparison of the predicted scores for knowledge entries 1, 3, and 5 demonstrates that the Value Model has successfully learned the concepts of ”maintaining a safe following distance” and ”respecting and yielding to pedestrians.”

In the second scenario (bottom, images 3 and 4), a large truck is present ahead of the ego vehicle. According to defensive driving principles, one should maintain a distance from large, risky vehicles. The comparison of predicted scores for knowledge entries 1 and 4 indicates that the Value Model has effectively learned the concepts of ”maintaining a safe following distance” and ”avoiding large vehicles.” In both scenarios, the scores predicted by our Value Model closely align with the GT.

The scores predicted by our value model in both scenarios align closely with the ground truth. This demonstrates the model’s ability to effectively gauge knowledge compliance across various contexts and assign reasonable scores accordingly.

## F. Sample of the Value Model Dataset

Some samples from the value model dataset are shown in Figure 8. We will make the entire dataset and the model publicly available in the future.


sents the ego vehicle. In the first case (top), the first trajec-

**Perception Results and Input Trajectories to the Value Model**

**Retrieved Knowledge Entries**

**Prediction**

**GT**



1. Yield to pedestrians and non-motor vehicles.
2. Do not honk to rush vulnerable pedestrians, including the elderly, disabled, or pregnant women.
3. Maintain a safe following distance at all times to prevent rear-end collisions.
4. Treat unexplained slow driving or frequent braking as a risk. Increase distance and overtake when safe.
5. Watch for sudden crossings, especially on roads without crosswalks. Cover the brake upon seeing pedestrians and be ready to stop.


0.7 1.0

0.4 0.0

1.0 1.0

0.9 0.9

0.9 0.9



1. Yield to pedestrians and non-motor vehicles.
2. Do not honk to rush vulnerable pedestrians, including the elderly, disabled, or pregnant women.
3. Maintain a safe following distance at all times to prevent rear-end collisions.
4. Treat unexplained slow driving or frequent braking as a risk. Increase distance and overtake when safe.
5. Watch for sudden crossings, especially on roads without crosswalks. Cover the brake upon seeing pedestrians and be ready to stop.


-0.5 -0.7

0.3 0.0

-0.7 -0.9

-0.3 -0.6

-0.5 -0.7



1. Large vehicles are heavy, stop slowly, and have blind spots. Always maintain a safety gap far exceeding that of regular cars.
2. Even with a green light, slow down and check the intersection. Be alert for sudden cross-traffic or pedestrians violating signals.
3. In traffic congestion, do not block the intersection.
4. Maintain a safe following distance sufficient for emergency braking.
5. Slow down and proceed with caution when pedestrians or non-motorized vehicles are present on the roadside.


0.4 0.8

0.8 1.0

0.6 0.9

0.8 1.0

0.9 1.0



1. Large vehicles are heavy, stop slowly, and have blind spots. Always maintain a safety gap far exceeding that of regular cars.
2. Even with a green light, slow down and check the intersection. Be alert for sudden cross-traffic or pedestrians violating signals.
3. In traffic congestion, do not block the intersection.
4. Maintain a safe following distance sufficient for emergency braking.
5. Slow down and proceed with caution when pedestrians or non-motorized vehicles are present on the roadside.

-0.7 -1.0

0.6 1.0

0.5 0.8

-0.6 -0.9

0.8 1.0

Figure 8. **Visualization of future state evaluation results using the proposed Value Model.** Displayed are BEV of the future state, retrieved knowledge entries, and Prediction vs. GT Scores. The comparison between compliant (Rows 1, 3) and risky (Rows 2, 4) trajectories demonstrates that our Value Model can evaluate future states against knowledge.