

Reviewer Rebuttal Report

Trust-Guided Multimodal LLM Integration with RL for Autonomous Driving

WACV 2026 LLVM-AD Workshop — Paper ID: LLVM-AD-13

Review Summary

Reviewer	Rating	Confidence	Decision
Program Chairs	–	–	Accept
Reviewer rm61	8/10	4/5	Clear Accept
Reviewer hNhx	8/10	4/5	Clear Accept

Overall: Strong Accept - Both reviewers rate the paper in the top 50% of accepted papers.

1 Concern 1: Trust Floor (τ_{min}) Not Ablated

“The lower bound ($\tau_{min} = 0.3$) is well motivated, but its sensitivity is not explored.”

Response: The response is included in **Section 7.3, “Trust Floor Sensitivity Analysis”**.

- **Content:** We conducted a systematic ablation across 5 values ($\tau_{min} \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$). Results in Table 7 demonstrate that 0.3 achieves optimal variance (2.1) and success rate (0.84).
- **Reflected in Manuscript:** Page 6, Section 7.3, Table 7.

2 Concern 2: Fallback Mechanisms Not Quantified

“While fallback mechanisms are mentioned, their quantitative impact is not evaluated.”

Response: The response is included in **Section 7.4, “Fallback Mechanism Quantification”**.

- **Content:** We quantified fallback rates across all scenarios. Table 8 reports that fallbacks handle 45% of decisions on average, reaching 91% in night driving, confirming robustness.
- **Reflected in Manuscript:** Page 7, Section 7.4, Table 8.

3 Concern 3: Reward Attribution

“Reward includes LLM-alignment term, which may partially explain gains.”

Response: The response is included in **Section 7.5, “Reward Component Attribution”**.

- **Content:** We ablated the R_{llm} term. Comparison shows that 11.8% of improvement comes from reward shaping, while 17.5% comes from trust-gated policy learning (total 29.3%), confirming the mechanism’s independent value.
- **Reflected in Manuscript:** Page 7, Section 7.5, Table 9.

4 Concern 4: No Comparison to Ensembles

“No comparison against alternative uncertainty estimation methods.”

Response: The response is included in **Section 7.1, Paragraph 3, “Comparison to Alternative Uncertainty Methods”**.

- **Content:** We added a dedicated discussion contrasting our approach with ensembles (too costly: $5 - 10\times$ RAM) and Bayesian methods (estimate observation uncertainty, not guidance reliability).
- **Reflected in Manuscript:** Page 6, Section 7.1.

5 Concern 5: Simulation-Only Evaluation

“Sim-to-real transfer remains unvalidated.”

Response: The response is included in **Section 8, “Limitations and Future Work”**.

- **Content:** We explicitly acknowledged this limitation. As an Algorithms Track submission, we validate the mechanism on a kinematic model with realistic parameters ($L = 2.7\text{m}$, 10Hz control).
- **Reflected in Manuscript:** Page 8, Section 8, Bullet 1.

6 Concern 6: High Computational Cost

“ ~ 70 ms inference latency.”

Response: The response is included in **Section 7.6, “Computational Cost”**.

- **Content:** We provided a detailed component breakdown (Table 10). The trust mechanism adds negligible overhead (2ms). Latency is dominated by LLM, which can be decoupled or quantized.
- **Reflected in Manuscript:** Page 7, Section 7.6, Table 10.

7 Concern 7: Model Specificity

“Trust mechanism trained with specific LLMs.”

Response: The response is included in **Section 8, “Limitations and Future Work”**.

- **Content:** We acknowledged this limitation. The principle of learning when to trust generalizes, though specific trust scores are model-dependent. Cross-model generalization is listed as future work.
- **Reflected in Manuscript:** Page 8, Section 8, Bullet 4.

Code-Paper Alignment

Concern	Paper Section	Code Location	Status
τ_{min} sensitivity	Section 7.3	run_ablations.py:48-53	Addressed
Fallback quant.	Section 7.4	RLAD1.py:2086-2118	Addressed
Reward attribution	Section 7.5	run_ablations.py:59	Addressed
Latency	Section 7.6	N/A	Addressed