
Task-Aligned Stability Analysis of Vision-Language Models for Autonomous Driving Hazard Detection

Anonymous Authors¹

Abstract

Vision-language models (VLMs) are increasingly used for scene understanding in autonomous driving, but robustness analysis often relies on task-agnostic embedding stability alone. We study whether corruption-induced embedding drift predicts changes in a task-aligned hazard score derived from CLIP image-text similarities. Using controlled corruptions on BDD100K road scenes, we compare embedding drift against margin drift, defined as the change in hazard score under perturbation. The relationship is highly corruption-dependent: some families exhibit strong coupling between representation drift and decision drift, while others induce hazardous decision instability despite relatively modest embedding change. A scatter analysis reveals a structured upper envelope relating margin drift to embedding drift, suggesting that representation change constrains but does not determine decision instability. These results suggest that robustness benchmarks should include task-aligned stability measures in addition to embedding-level perturbation statistics.

1. Introduction

Autonomous vehicles depend on perception systems that remain reliable under distribution shift. Hazard detection is a particularly safety-critical instance of this problem because small visual changes can alter the driving decision implied by a scene. Vision-language models (VLMs) are attractive for this setting because they combine visual recognition with prompt-based semantic reasoning, enabling natural-language descriptions of hazardous and non-hazardous traffic scenarios (Bordes et al., 2024; Radford et al., 2021). However, prior work shows that VLM-based hazard detectors can be brittle, overreact in ambiguous settings, and impose substantial computational cost (Choi et al., 2025; Adil et al., 2025).

A common robustness perspective is to measure how much the image embedding changes under corruption. While useful, this quantity is task-agnostic: a large change in

embedding space need not produce a large change in the decision variable that matters for safety. We therefore contrast *embedding drift* with a task-aligned *margin drift* based on a CLIP hazard score. This aligns with broader efforts toward quantitative reliability analysis, robustness under distribution shift, and principled benchmark design.

Our central question is simple: *to what extent does corruption-induced embedding drift predict decision-relevant instability in VLM hazard detection?* We find that the answer depends strongly on corruption family. JPEG compression and downsampling produce relatively predictable behavior, whereas motion blur and occlusion can trigger high decision flip rates without proportionally large embedding movement. This suggests that representation stability alone is an incomplete proxy for safety-relevant robustness.

2. Method

Hazard scoring. Let $\hat{f}(x)$ and $\hat{g}(t)$ denote frozen CLIP image and text encoders. We define hazardous prompts H and non-hazardous prompts N , then score image x by

$$\Lambda(x) = \max_{h \in H} \langle \hat{f}(x), \hat{g}(h) \rangle - \max_{n \in N} \langle \hat{f}(x), \hat{g}(n) \rangle, \quad (1)$$

where larger values indicate greater hazard likelihood. Positive scores correspond to hazardous scenes and negative scores to non-hazardous scenes.

Our prompt sets are intentionally simple and interpretable. Hazardous prompts include phrases such as *a pedestrian crossing the road*, *a cyclist in the lane*, and *a stopped vehicle blocking traffic*. Non-hazardous prompts include *a clear road with no obstacles*, *a normal highway scene*, and *a clear lane ahead*. This follows the prompt-based hazard scoring strategy used in recent autonomous-driving work built on CLIP (Radford et al., 2021; Greer et al., 2026).

Corruptions and drift metrics. For corruption family c and severity $\alpha \in \{1, \dots, 5\}$, let $T_{c,\alpha}(x)$ denote the corrupted image. We evaluate seven families: fog, Gaussian blur, motion blur, JPEG compression, low light, occlusion, and downsampling. We then define

$$\Delta(x, c, \alpha) = \|\hat{f}(x) - \hat{f}(T_{c,\alpha}(x))\|_2 \quad (2)$$

Algorithm 1 Task-aligned corruption evaluation

```

Require: Images  $\mathcal{X}$ , corruptions  $\mathcal{C}$ , severities  $\mathcal{A}$ , prompt
collections  $H, N$ 
0: for  $c \in \mathcal{C}$  do
0:   for  $\alpha \in \mathcal{A}$  do
0:     for  $x \in \mathcal{X}$  do
0:       Compute  $\Lambda(x)$  and  $\hat{f}(x)$ 
0:       Form  $x' = T_{c,\alpha}(x)$ 
0:       Compute  $\Lambda(x')$  and  $\hat{f}(x')$ 
0:       Record  $\Delta = \|\hat{f}(x) - \hat{f}(x')\|_2$ 
0:       Record  $D_\Lambda = |\Lambda(x) - \Lambda(x')|$ 
0:       Record  $\text{flip} = \mathbb{1}[\text{sign}(\Lambda(x)) \neq$ 
 $\text{sign}(\Lambda(x'))]$ 
0:     end for
0:   end for
0: end for

```

and

$$D_\Lambda(x, c, \alpha) = |\Lambda(x) - \Lambda(T_{c,\alpha}(x))|. \quad (3)$$

Here Δ is task-agnostic embedding drift and D_Λ is task-aligned margin drift.

We also record a binary *flip* indicator,

$$\text{flip} = \mathbb{1}[\text{sign}(\Lambda(x)) \neq \text{sign}(\Lambda(T_{c,\alpha}(x)))], \quad (4)$$

which captures whether corruption changes the hazard verdict.

3. Experimental setup

We use 2,000 images sampled from the BDD100K validation split (Yu et al., 2020). For each image, corruption family, and severity, we compute the metrics above and then summarize each family using: (i) Spearman correlation $\rho(\Delta, D_\Lambda)$, (ii) flip rate, and (iii) mean drift as a function of severity. This design isolates the relationship between representation change and decision instability under controlled shift.

4. Results

4.1. Hazard Score Behavior

The CLIP-based score produces plausible rankings: low-scoring scenes are typically open roads with minimal interaction, while high-scoring scenes contain dense traffic, nearby actors, or explicit obstructions. This supports the use of Λ as a continuous decision variable for analyzing corruption effects, even in the absence of a fully supervised detection benchmark.

Table 1. Family-wise coupling between embedding drift and decision drift.

Corruption family	$\rho(\Delta, D_\Lambda)$	Flip rate
Fog	0.348	0.108
Gaussian blur	0.587	0.162
Motion blur	0.462	0.242
JPEG compression	0.738	0.108
Low light	0.553	0.154
Occlusion	0.297	0.178
Downsampling	0.615	0.059
Average	0.514	0.144

4.2. Embedding Drift and Margin Drift Are Not Interchangeable

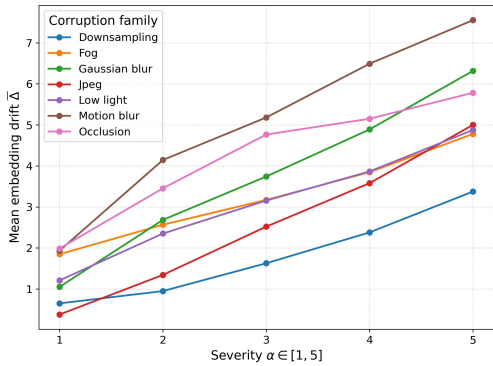
Table 1 shows strong heterogeneity across corruption families. JPEG compression and downsampling have relatively strong rank correlation between embedding drift and margin drift, suggesting that representation movement is a reasonable proxy for decision movement in those cases. In contrast, occlusion and fog show much weaker coupling. Most notably, motion blur exhibits the highest flip rate even though its drift correlation is only moderate. Thus, the corruption families that most threaten decision reliability are not necessarily those that induce the largest embedding perturbations.

Figure 1 summarizes these patterns visually. The left panel plots mean embedding drift against severity, and the right panel plots mean margin drift against severity. Both generally increase monotonically, confirming that stronger corruptions perturb both representation and task score. However, the relative ordering differs across families. Motion blur and occlusion induce disproportionately large margin effects relative to their embedding movement, while JPEG compression causes substantial embedding change with comparatively lower decision instability.

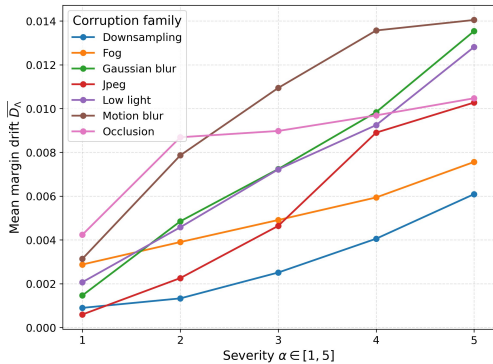
4.3. Scatter Geometry of Embedding and Margin Drift

To further analyze the relationship between representation change and decision instability, we examine the joint distribution of embedding drift Δ and margin drift D_Λ across corruption families. Figure 2 visualizes this relationship. Several consistent structural patterns emerge:

Upper-envelope (conical) constraint. Across all corruption families (Fig. 2a), the data occupies a cone-shaped region with a clear upper boundary relating D_Λ to Δ . Large margin shifts are only observed when embedding drift is non-negligible, suggesting that representation change is a necessary condition for decision change. However, the spread within this cone is substantial, indicating that embedding drift alone does not determine the magnitude of decision impact.



(a) Mean embedding drift vs. severity.



(b) Mean margin drift vs. severity.

Figure 1. Severity trends differ across corruption families. Similar monotonicity does not imply similar safety impact.

Family-dependent coupling regimes. Different corruption families populate distinct regions of this space. JPEG compression (Fig. 2b) exhibits a tight, approximately linear relationship between Δ and D_Λ , consistent with its high Spearman correlation in Table 1. In contrast, motion blur and occlusion (Fig. 2c–d) display much broader vertical dispersion. For these families, similar levels of embedding drift can correspond to widely varying margin shifts, indicating weaker coupling between representation and decision.

Low-drift, high-impact regime. Of particular concern is the region characterized by relatively small Δ but large D_Λ . This regime appears prominently in motion blur and occlusion. Points in this region correspond to perturbations that minimally affect the global embedding while significantly altering the relative similarity to hazardous versus non-hazardous prompts. Empirically, these cases account for a large fraction of observed flips, suggesting that they represent a critical failure mode for VLM-based hazard detection.

Severity stratification. Coloring by severity reveals that embedding drift increases relatively smoothly with corruption strength, while margin drift exhibits more heteroge-

neous growth. In tightly coupled families (e.g., JPEG), increasing severity moves points outward along a narrow band. In weakly coupled families (e.g., occlusion), severity increases both the mean and the variance of margin drift, further amplifying unpredictability at higher corruption levels.

Interpretation. The scatter structure indicates that embedding drift constrains but does not determine decision change. While large margin shifts require non-negligible representation movement, similar levels of embedding drift can lead to substantially different margin outcomes depending on corruption type. This asymmetry highlights the role of corruption-specific effects in shaping decision stability.

Hypothesis: bounded decision sensitivity. The observed upper-envelope structure suggests a potential Lipschitz-like relationship between embedding drift and margin drift of the form $D_\Lambda \leq L \cdot \Delta$ for some task-dependent constant L . While we do not formally estimate or validate such a bound in this work, the empirical geometry indicates that decision changes are constrained by representation movement but not tightly determined by it. Future work could aim to characterize this constant analytically or empirically, potentially yielding formal guarantees on decision stability under bounded embedding perturbations.

4.4. Implications for Benchmarking

These results suggest that robustness evaluation for multimodal systems should incorporate task-aligned metrics in addition to representation-level measures. In particular, reporting margin drift and threshold-flip rates by corruption family provides a more direct view of decision stability under shift. Benchmarks that rely solely on embedding perturbation statistics may overlook failure modes that have limited representation impact but significant decision consequences.

5. Conclusion

We presented a compact, task-aligned stability analysis for CLIP-based hazard detection under controlled image corruption. Across seven corruption families, embedding drift and hazard-margin drift were related but far from equivalent. Some families produced predictable coupling between the two, while others created dangerous decision instability that was not well reflected by embedding movement alone. For theory-informed benchmark design, this suggests that representation robustness should be complemented by task-aligned stability metrics that quantify how corruption propagates into decision variables and final verdicts.

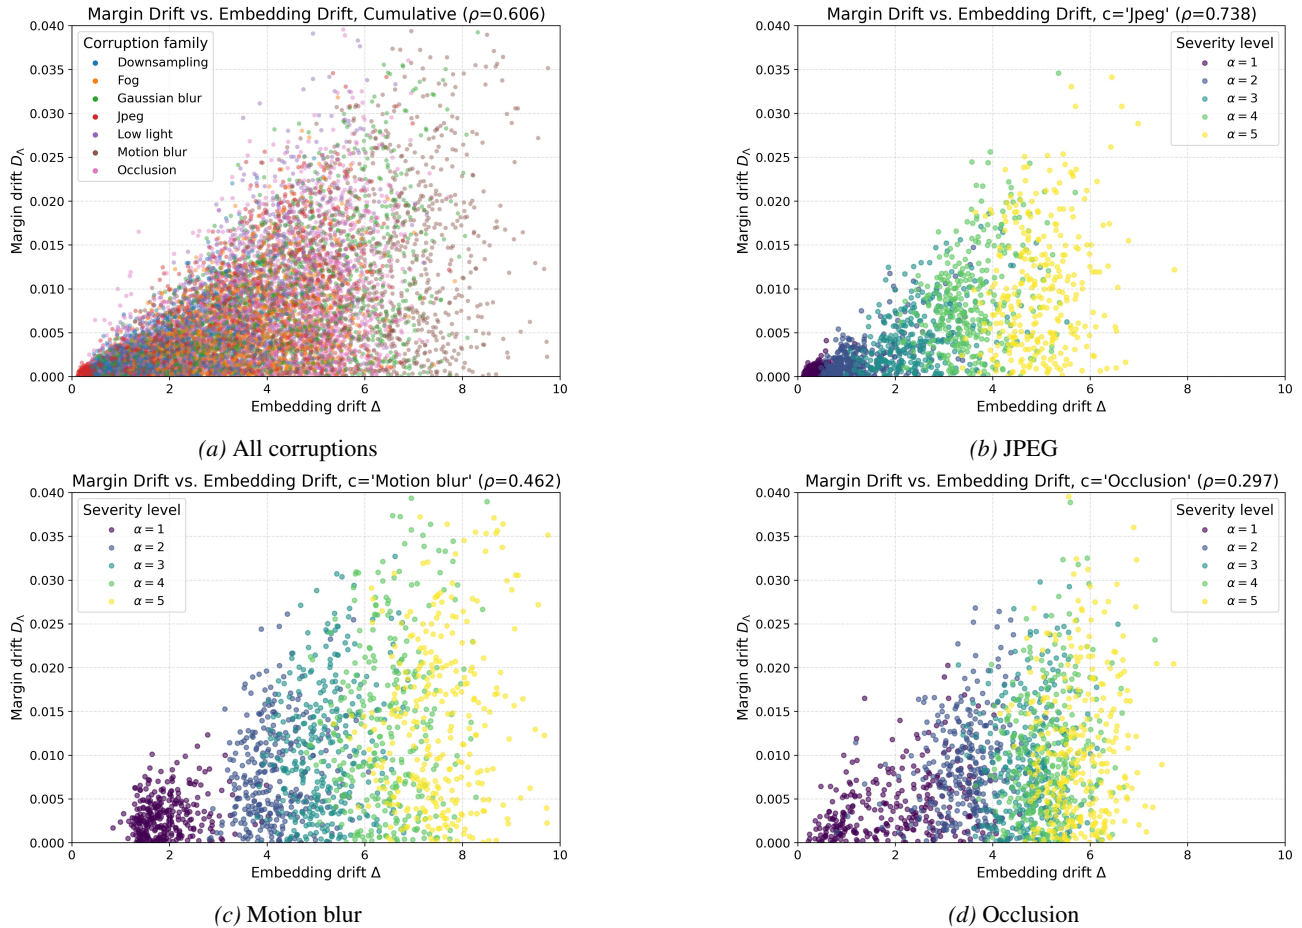


Figure 2. Embedding drift vs. margin drift across corruption families. Points are colored by corruption family in (a) and by severity level in (b)–(d).

References

- Adil, M., Lee, G., Gonzalez, V. A., and Mei, Q. Using vision language models for safety hazard identification in construction, 2025. URL <https://arxiv.org/abs/2504.09083>.
- Bordes, F., Pang, R. Y., Ajay, A., Li, A. C., Bardes, A., Petryk, S., Mañas, O., Lin, Z., Mahmoud, A., Jayaraman, B., Ibrahim, M., Hall, M., Xiong, Y., Lebensold, J., Ross, C., Jayakumar, S., Guo, C., Bouchacourt, D., Al-Tahan, H., Padthe, K., Sharma, V., Xu, H., Tan, X. E., Richards, M., Lavoie, S., Astolfi, P., Hemmat, R. A., Chen, J., Tirumala, K., Assouel, R., Moayeri, M., Talattof, A., Chaudhuri, K., Liu, Z., Chen, X., Garrido, Q., Ullrich, K., Agrawal, A., Saenko, K., Celikyilmaz, A., and Chandra, V. An introduction to vision-language modeling, 2024. URL <https://arxiv.org/abs/2405.17247>.
- Choi, D., Lee, S., and Song, Y. Better safe than sorry? overreaction problem of vision language models in visual emergency recognition, 2025. URL <https://arxiv.org/abs/2505.15367>.
- Greer, R., Keskar, M., Martinez-Sanchez, A., Roy, P., Shriram, S., and Trivedi, M. Vision and language: Novel representations and artificial intelligence for driving scene safety assessment and autonomous vehicle planning, 2026. URL <https://arxiv.org/abs/2602.07680>.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., and Sutskever, I. Learning transferable visual models from natural language supervision, 2021. URL <https://arxiv.org/abs/2103.00020>.
- Yu, F., Chen, H., Wang, X., Xian, W., Chen, Y., Liu, F., Madhavan, V., and Darrell, T. Bdd100k: A diverse driving dataset for heterogeneous multitask learning, 2020. URL <https://arxiv.org/abs/1805.04687>.

A. Prompt sets and corruption ranges

Hazardous prompts: *a pedestrian crossing the road, a cyclist in the lane, a stopped vehicle blocking traffic, road debris on the street, a construction zone on the road, an emergency vehicle with flashing lights*. Non-hazardous prompts: *a clear road with no obstacles, a normal highway scene, an empty intersection, a clear lane ahead, a normal driving scene*. Corruption severities are uniformly discretized over the same ranges used in the class draft: fog density [0.15, 0.75], Gaussian blur $\sigma \in [1, 8]$, motion-blur kernel size [3, 25], JPEG quality [90, 10], low-light scale [0.8, 0.2], occlusion mask ratio [0.05, 0.60], and downsampling scale [0.90, 0.20].