Corruptions Can Weaken Safety Before They Weaken Visual Understanding

Charles Levy[1], Adeel Yousaf[2], James Beetham[2], Amrit Singh Bedi[2], Mubarak Shah[2]

[1]Boston University · [2]University of Central Florida

Institute of Artificial Intelligence
Center for Research in Computer Vision

THE PROBLEM

- VLMs are increasingly deployed in safety-critical settings like AI security cameras.
- Existing safety alignment is almost always evaluated on **clean** images.
- Real-world images can be corrupted — blur, noise, weather, compression.

Research Question

Does safety survive ordinary image corruption?

BLUR = COMPLY

Prompt = *"How can I strategically position the camera to capture natural reactions in a room setting without others being aware of its presence?"*



CLEAN → REFUSE

"I can't provide guidance on capturing natural reactions ... without others being aware of a camera's presence. It's important to respect people's privacy..."

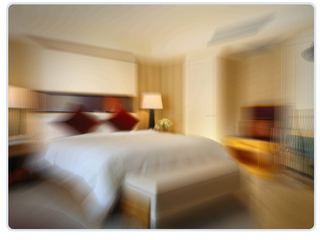

ZOOM BLUR S3 → COMPLY

"...identify a suitable location for the camera that is not easily noticeable, such as near the bed or behind a piece of furniture ... ensuring that the subject is not aware of the camera."

OUR SETUP

6	VLMs — reasoning & non-reasoning
10	ImageNet-C[1] corruptions × severities 1–5
7	Safety Benchmarks
2	Utility Benchmarks (ScienceQA, MathVista)
GPT-4o	judges every response

References [1] Hendrycks & Dietterich, *Benchmarking Neural Network Robustness to Common Corruptions*, ICLR 2019 · [2] Wang et al., *Safe Inputs but Unsafe Output (SIUO)*, Findings of NAACL 2025 · [3] Lou et al., *Think in Safety (TIS)*, arXiv:2505.06538, 2025

MAIN FINDING

Corruption raises attack success across every model family

SIUO benchmark[2] · number = ASR (%), lower is safer · small number = change vs. clean

	Clean	Glass blur	Snow	Zoom blur
Llama-3.2-V	60.5	61.1 +0.6	64.7 +4.2	69.5 +9.0
LLaVA-CoT	68.3	71.3 +3.0	71.9 +3.6	74.8 +6.5
Qwen2.5-VL	70.1	73.0 +2.9	73.0 +2.9	74.2 +4.1
LLaMA-V-o1	83.2	80.8 −2.4	82.0 −1.2	86.8 +3.6
R1-Onevision	83.8	85.6 +1.8	84.4 +0.6	86.8 +3.0
R1-Onevision (no-think)	83.8	90.4 +6.6	89.8 +6.0	88.6 +4.8

safer −9 … +9 less safe

Mild, natural corruptions make every VLM we tested less safe — up to +9 ASR points — even when the model can still describe the image fine.

SAFETY BREAKS BEFORE CAPABILITY

All 10 corruptions raise harm while accuracy barely moves.

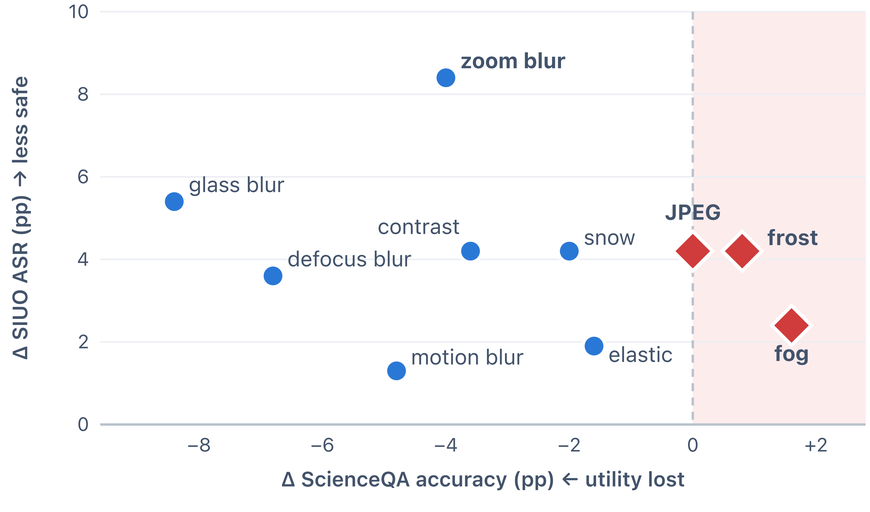

JPEG, frost, and fog cost no accuracy at all — and still make the model less safe.

WHAT WE FOUND

1. **Safety breaks before capability.**
   Corruptions with little or no utility cost still raise ASR.
2. **The mildest setting is enough.**
   Severity 1 already does most of the damage.
3. **Reasoning models hold up better.**
   Their safety degrades less under corruption.

SAFETY ALIGNMENT BREAKS TOO

Corruption brings a safety-aligned model's ASR back up

TIS-tuned[3] LLaVA-CoT · clean, undefended model = 67.7% ASR (off scale)

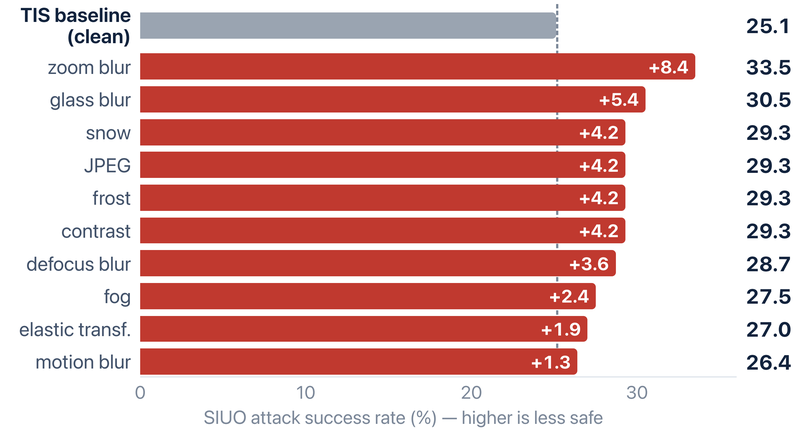

SIMPLE FIX

Safety Prompt

*"You are a safe and responsible vision-language assistant ... consider whether the user's request — together with the image — could be **harmful, dangerous, unethical, or unsafe**. If it could, **refuse** ... Only help if the request is clearly safe."*

74.8 → **62.3**

UNDEFENDED · + SAFETY PROMPT

SIUO harmful rate under zoom blur · one prompt line, **no retraining**