Driftline Research Note 001

Troy Pressley
Driftline Research (Independent)
Houston, Texas — March 2026

Structural Stability Failures in Diffusion Image Models
An observational research note

In a Driftline baseline generation run, approximately 7,000 images were produced using a minimal prompt workflow. From this run, 280 outputs were identified as chair-class images. A randomly sampled subset of 58 chair images was manually reviewed and scored using the Structural Validity Score (SVS), a four-level rubric designed to test whether diffusion-generated objects preserve the mechanical logic of rigid physical structures. Within this 58 image subset, 16 images were classified as SVS-3 (structurally plausible) and 13 as SVS-2 (minor irregularities but mechanically usable). The remaining images exhibited structural failure: 14 were classified as SVS-1 (major structural issues) and 15 as SVS-0 (complete structural collapse). This produces an even split in the scored subset: 29 structurally acceptable outputs and 29 structurally failed outputs.

Method note. Images were generated using Stable Diffusion XL through the ComfyUI inference framework under minimal prompt conditions (typically the prompt "chair", occasionally with neutral background or studio wording). Random seeds were used across runs with consistent resolution settings. From approximately 7,000 generated images, 280 were identified as chair-class outputs through manual review of the baseline generation batch. A randomly sampled subset of 58 images was then scored using the Structural Validity Score (SVS). The baseline configuration used a 768 x 1536 latent image size with approximately 30 sampling steps.

Structural Validity Score (SVS)

SVS-3 - Structurally valid. The object maintains coherent load paths and could plausibly function as a real chair.
SVS-2 - Minor structural anomalies. The chair is mostly plausible but contains small distortions or non-critical structural irregularities.
SVS-1 - Major structural failure. Key structural members are malformed, disconnected, or mechanically unstable.
SVS-0 - Structural collapse. The object cannot plausibly exist as a functional physical structure.

Figure 1 - Structurally valid chair (SVS-3). The seat, back, and supporting legs form a coherent load path. Although stylistically imperfect, the object remains mechanically plausible.

Failure A (SVS-0). Duplicated seat planes create impossible stacked geometry caused by hallucinated pixels.

Failure B (SVS-1). Structural members deform and disconnect, breaking the mechanical integrity of the frame.

Failure C (SVS-0). The office chair form collapses into an unsupported secondary seat structure.

Interpretation. These examples illustrate a recurring pattern observed during large-batch diffusion generation: models frequently preserve the surface appearance of objects while violating their underlying structural logic. The model is not reasoning about joinery, center of mass, or load-bearing geometry. Instead it generates pixels that statistically resemble chairs while occasionally producing

visually convincing but mechanically impossible objects. **Within this exploratory subset, roughly half of generated chairs fail a basic structural plausibility test.**

Limitations. This document reports an exploratory observational baseline rather than a formal statistical evaluation. The SVS rubric represents a human judgment framework applied to a small subset of generated images and has not yet been validated through inter-rater reliability testing. Results are therefore illustrative and intended to document a recurring failure pattern rather than establish a generalizable benchmark for diffusion models.