Probing Visual Idiosyncrasies in Text-to-Image Models

Yongqian Peng* **Hongjie Li*** **Yichao Zhou*** **Chi Zhang†** **Yixin Zhu†**

Peking University

chizhang.cz@pku.edu.cn, yixin.zhu@pku.edu.cn

Abstract

Text-to-Image (T2I) models turn short text descriptions into images, and the leading models now produce images of comparable visual quality. Whether their outputs still carry model-specific characteristics matters for attribution, for model comparison beyond aggregate quality, and for evaluations that assume anonymity. We ask whether images from frontier T2I models can still be classified by source, and which visual properties make this possible. Here we show that a standard ResNet-50 identifies the source of an image among HiDream-O1-Dev, FLUX.2-dev, HunyuanImage-3.0, and Qwen-Image-2512, the four highest-ranked open-source models on the Text-to-Image Arena, with 100.00% accuracy on held-out prompts. Human participants reach only 62.61% on the pairwise version of this task, so the distinguishing cues are not visually obvious; they are also not tied to a particular classifier and survive severe image corruptions and out-of-distribution prompts. We trace them to *visual idiosyncrasies* at three levels: low-level texture statistics alone let a lightweight multi-layer perceptron match ResNet-50; object layout, depth, and albedo estimated from the images each carry a source signal; and image captions reveal model-specific preferences in style and background that the prompts never specify. A one-sentence hint about these cues raises human accuracy by 9.03 percentage points. These idiosyncrasies make source classification a tool for measuring similarity between models and between real and generated images, and they caution against assuming anonymity in human preference evaluation of T2I models. Project page: https://ppyyqq.github.io/whichmodel/.

Introduction

As Text-to-Image (T2I) models rapidly advance and produce images of increasingly comparable visual quality (Arena 2026), whether their outputs still carry model-specific characteristics becomes a pressing question. Prior work has shown that generated images can be attributed to their source models (Girish et al. 2021; Bui, Yu, and Collomosse 2022; Corvi et al. 2023b; Guarnera, Giudice, and Battiato 2024), and most recently that frontier T2I sources remain distinguishable across models and generation settings (Xu, Zhang, and Shi 2025). These studies, however, treat attribution as the end goal and leave open what makes the sources distinguishable: the visual properties that systematically characterize each model remain largely unexplored. Characterizing these properties reveals how models differ beyond aggregate quality, supports model comparison and real-synthetic similarity analysis, and bears on evaluation protocols that assume model identities are hidden.

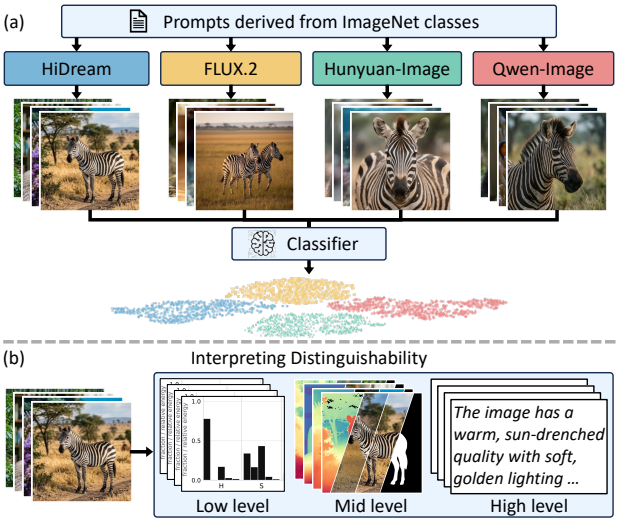

Figure 1: **Source classification and multi-level cue analysis.** (a) Images generated by four T2I models from shared ImageNet-derived prompts are highly distinguishable by source under a trained classifier. (b) We further investigate this distinguishability by analyzing potential model-specific cues at low, mid, and high visual levels.

In this work, we ask whether outputs from frontier T2I models can still be reliably classified by source, and which recurring visual properties enable this distinction; we call such properties *visual idiosyncrasies*. We adopt a controlled source-classification task in the spirit of the "Name That Dataset" analyses of dataset bias (Torralba and Efros 2011; Liu and He 2025) and of recent work on idiosyncrasies of Large Language Models (LLMs) (Sun et al. 2025): every model receives the same prompts, and a classifier must predict the source model from the image alone. We instantiate the task with four frontier open-source models (HiDream-O1-Dev (HiDream), FLUX.2-dev (FLUX.2), HunyuanImage-3.0 (HunyuanImage), and Qwen-Image-2512 (Qwen-Image)) and dissect the cues that make

*These authors contributed equally.

†Corresponding author.

their outputs distinguishable at the low, mid, and high levels of the visual hierarchy (Fig. 1).

A standard ResNet-50 (He et al. 2016) identifies the source with 100.00% accuracy on prompts held out from training, far above the 25% chance level. This distinguishability is not a matter of visually obvious differences: on the pairwise version of the task, human participants average 62.61%, whereas ResNet-50 averages 99.99%. It is also not tied to a particular classifier, persists under severe image corruptions, and generalizes to out-of-distribution (OOD) prompts.

We then characterize the visual idiosyncrasies behind this result at three levels. At the *low level*, a bank of 225 handcrafted features suffices to distinguish the models, and a category-wise analysis singles out its 38 texture features: a lightweight Multi-Layer Perceptron (MLP) trained on texture alone reaches 99.76% four-way accuracy and generalizes to new prompts at least as well as ResNet-50. At the *mid level*, object layout, depth, and albedo estimated from the images each carry a source signal, from 59.50% accuracy for layout to 99.75% for albedo, and the models differ systematically in average object shape, depth range, focal length, and the spatial scale of albedo variation. At the *high level*, image captions expose preferences that the prompts never specify: using only Term Frequency-Inverse Document Frequency (TF-IDF) features of the captions, logistic regression classifiers reach up to 68.78% and 71.78% four-way accuracy from style and background descriptions, respectively.

These idiosyncrasies are useful beyond explaining the classifier. A one-sentence hint describing them raises human accuracy from 62.61% to 71.64%. Source classification also doubles as a similarity measure: leave-one-model-out prediction reveals which models resemble one another, and passing real ImageNet images through the classifiers at each level shows which model real photographs most resemble. Finally, persistent idiosyncrasies can betray a model's identity, which cautions against assuming anonymity in human preference evaluation such as the Text-to-Image Arena. Together, these results establish source classification as an interpretable probe of how frontier T2I models differ, and visual idiosyncrasies as the properties it detects.

Related Work

Classification-based analysis of visual datasets. Torralba and Efros (2011) introduce the "Name That Dataset" experiment, in which a classifier that predicts which dataset an image comes from exposes dataset-specific bias. Such bias persists in modern large-scale datasets (Liu and He 2025) and arises from both structural and semantic factors (Zeng, Yin, and Liu 2024). Our study is motivated by Liu and He (2025) and adopts the classification-based analyses of Liu and He (2025) and Zeng, Yin, and Liu (2024), but targets model-specific rather than dataset-specific patterns: the classes are the source models of generated images from the same prompts, and the object of study is the visual properties that make them distinguishable.

Detection and attribution of AI-generated images. Synthetic images contain model-specific artifacts that enable detection (Marra et al. 2019; Chai et al. 2020; Wang et al. 2020); these artifacts take the form of spatial patterns, frequency signatures, textures, and style cues (Liu, Qi, and Torr 2020; Frank et al. 2020; Corvi et al. 2023a; Amoroso et al. 2024). Recent detectors extend to diffusion models and generalize better with pretrained features (Wang et al. 2023a; Ojha, Li, and Lee 2023; Cozzolino et al. 2024a; Liu et al. 2024b), and zero-shot, training-free detectors exploit compression likelihood, autoencoder reconstruction, perturbation sensitivity, and diffusion-model geometry or dynamics (Cozzolino et al. 2024b; Ricker, Lukovnikov, and Fischer 2024; Brokman et al. 2025; Choi, Lee, and Lee 2025; Liang et al. 2025b; Choi et al. 2026). Attribution asks not whether an image is synthetic but which model produced it. Yu, Davis, and Fritz (2019) show that GANs carry distinctive fingerprints, motivating attribution methods for GANs (Girish et al. 2021; Bui, Yu, and Collomosse 2022) and for diffusion and other generative models (Wang et al. 2023b; Corvi et al. 2023b; Sha et al. 2023; Wang et al. 2024; Laszkiewicz et al. 2024; Liu et al. 2024a; Nguyen, Azizpour, and Stamm 2025; Wang et al. 2026). Most closely related, Xu, Zhang, and Shi (2025) study source distinguishability of T2I models under varying models and settings, and Naseh et al. (2026) deanonymize arena leaderboards through prompt-conditioned image-embedding clusters. Both establish that the source is recoverable; we ask what makes it recoverable and characterize the underlying visual idiosyncrasies at three levels.

Characterizing generative model idiosyncrasies. Beyond artifacts for detection and attribution, recent work characterizes properties of the generative models themselves. Sarkar et al. (2024) reveal systematic geometric inconsistencies in generated images, and Lee, Koo, and Kwak (2026) show that underspecified prompts expose model-specific semantic preferences. In the language domain, Sun et al. (2025) use controlled classification to identify idiosyncrasies across frontier LLMs. We bring this classification-based lens to images and characterize the idiosyncrasies of frontier T2I models across low-, mid-, and high-level representations.

Probing Visual Distinguishability of T2I Models

Formulation. We formulate source classification as a controlled task. Let $\mathcal{M} = \{m_1, \ldots, m_N\}$ denote a set of T2I models and $\mathcal{P}$ a shared prompt set. For each prompt $p \in \mathcal{P}$ and model $m_i \in \mathcal{M}$, we sample an image $x \sim m_i(\cdot \mid p)$ and assign it the source label $y = i$. A classifier observes only x and predicts y. Because every model receives the same prompts, prompt content cannot explain above-chance accuracy; accuracy above the balanced chance level of $1/N$ therefore indicates visual properties that consistently characterize a model across prompts, which we term *visual idiosyncrasies*.

Experimental setup. We instantiate the task with four recent T2I models: HiDream (Cai et al. 2026), FLUX.2 (Black Forest Labs 2025), HunyuanImage (Cao et al. 2025), and Qwen-Image (Qwen Team 2025). As of June 2026, they are the four highest-ranked open-source models on the Text-to-Image Arena (Arena 2026), each in the overall top 40.

HiDream	FLUX.2	HunyuanImage	Qwen-Image	Acc.
✓	✓			100.00
✓		✓		100.00
✓			✓	99.94
	✓	✓		100.00
	✓		✓	100.00
		✓	✓	100.00
✓	✓	✓	✓	100.00

Table 1: **Source classification accuracy across model combinations.** ResNet-50 accuracy (%) on the held-out ImageNet-derived test set for all six pairwise tasks and the four-way task. Every combination is nearly perfectly distinguishable.

Classifier	Acc.	Acc. (Random labels)
ResNet-50	100.00	26.41
Two-layer CNN	98.81	26.63
DINOv2 ViT-B/14	88.54	25.94
CLIP ViT-B/32	90.77	24.97

Table 2: **Source classification across classifier architectures.** Four-way accuracy (%) on the held-out ImageNet-derived test set, and accuracy after training on randomly permuted labels. All classifiers are far above the 25% chance level, and random-label performance remains near chance.

For each model, we generate images from prompts derived from the 1,000 ImageNet classes (Russakovsky et al. 2015), with four random seeds per prompt, yielding 4,000 images per model and 16,000 in total. Images are generated at 1024×1024 with each model's default sampling configuration (examples in Fig. A10). We hold out 200 classes as the in-distribution (ID) test set and the remaining classes as training classes. Our default classifier is a ResNet-50 (He et al. 2016) trained from random initialization on images resized to 224×224. Prompt construction and training details are in Secs. A.1 and A.2.

High Distinguishability Across T2I Models

Consistency across source combinations and classifiers. Table 1 reports ResNet-50 accuracy for all pairwise tasks and the four-way task. Every pair is perfectly distinguishable; the hardest, HiDream versus Qwen-Image, still reaches 99.94%. Expanding the label space to all models barely matters: the four-way classifier reaches 100.00%. This distinguishability is not specific to one architecture. A two-layer Convolutional Neural Network (CNN) trained from scratch and linear classifiers on frozen DINOv2 ViT-B/14 (Oquab et al. 2024) and CLIP ViT-B/32 (Radford et al. 2021) features differ widely in capacity, inductive bias, and pretraining objective, yet all remain far above chance (Tab. 2; confusion matrices in Sec. B.1). Nor does it stem from arbitrary correlations in the training labels: with randomly permuted labels, every classifier drops to the 25% chance level (Tab. 2, right column).

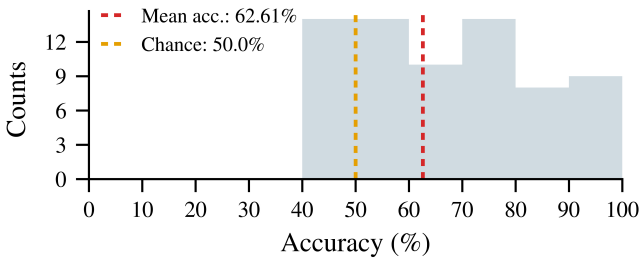

Figure 2: **Human accuracy on pairwise source classification.** Participants ($n = 69$) average 62.61% with wide individual variation, far below the 99.99% that ResNet-50 averages on the same model pairs.

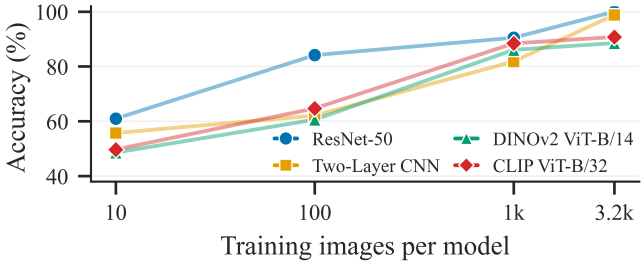

Figure 3: **Accuracy as a function of training set size.** Accuracy on the held-out test set rises steadily with the number of training images per model, and 10 images per model already give non-trivial accuracy.

Model differences are not fully perceptible to humans. To test whether this distinguishability rests on visually obvious differences, we conduct a two-alternative forced-choice study on Prolific. Each of 69 participants is assigned to one of the six model pairs, studies a reference board of nine images per model labeled only as Set A and Set B, and then classifies ten new images by source (Sec. A.3). Participants average 62.61% accuracy, far below the 99.99% that ResNet-50 averages over the same pairwise tasks, and vary widely (Fig. 2). Visually obvious differences alone therefore cannot account for the distinguishability: the classifiers capture model-specific patterns more consistently than humans do.

Distinguishability Persists Across Conditions

Accuracy improves with more training data. We vary the number of training images and retrain all classifiers (Fig. 3). Accuracy on the held-out test set rises steadily with training size, and even 10 images per model give non-trivial accuracy (*e.g.*, 60.96% for ResNet-50), well above chance.

Robustness to image corruptions. To test robustness to input degradation, we next train and evaluate ResNet-50 on images subjected to one of four corruptions: color jitter (Krizhevsky, Sutskever, and Hinton 2012), Gaussian blur, blending with Gaussian noise, and resolution reduction. Classification stays nearly perfect under moderate settings and degrades only under severe ones, yet remains well above chance even at a noise blending weight of 0.8 or a resolution of 2×2 (Tab. 3; examples in Fig. A11). Two further trans-

Corruption	Setting	Acc.
None	-	100.00
Color jittering	Strength: 1.0	100.00
Color jittering	Strength: 2.0	99.94
Gaussian blur	Radius: 4.0	100.00
Gaussian blur	Radius: 8.0	100.00
Gaussian noise	Blending weight: 0.4	99.91
Gaussian noise	Blending weight: 0.8	97.99
Low resolution	Resolution: 32×32	100.00
Low resolution	Resolution: 8×8	99.21
Low resolution	Resolution: 2×2	84.74

Table 3: **Robustness to image corruptions.** ResNet-50 accuracy (%) when training and test images are subjected to the same corruption. Accuracy stays nearly perfect under moderate settings and well above chance under the most severe ones; the full sweeps are in Tab. A3.

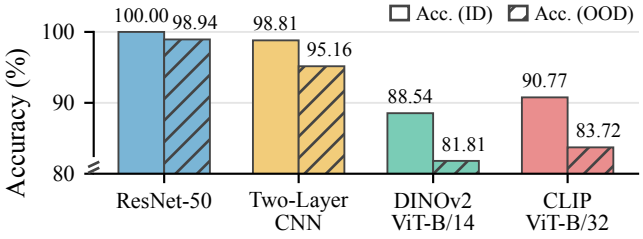

Figure 4: **Cross-prompt generalization.** Classifiers are trained on images generated from ImageNet-derived prompts and evaluated on images generated from Winoground-derived prompts (OOD). All four remain far above the 25% chance level.

formations show that the signal is neither purely global nor a fragile decoder artifact: shuffling 32×32 patches leaves accuracy at 99.43%, and re-encoding every image through an off-the-shelf Variational Autoencoder (VAE) leaves it above 99% (Sec. B.3).

Generalization to out-of-distribution prompts. Finally, we train on ImageNet-derived prompts and evaluate on 1,600 images per model generated from the 400 Winoground prompts (Thrush et al. 2022), which describe compositional relations among multiple entities rather than single object categories. All classifiers remain far above chance, and ResNet-50 retains 98.94% (Fig. 4). The distinguishability thus extends beyond the training prompt distribution to a different kind of visual content.

Characterizing Visual Idiosyncrasies of T2I Models

Having established that classifiers identify the source model reliably, we ask which model-specific properties make this possible. We analyze three complementary levels: low-level visual cues, mid-level visual representations, and high-level semantic realization. Qualitative examples for all three analyses are collected in Sec. F.

Feature category	Acc. (Single)		Acc. (Perm.)	
	ID	OOD	ID	OOD
Full	99.84	98.55	-	-
Tone	91.27	85.18	98.85	94.88
Color	93.28	91.10	96.36	92.93
Texture	**99.76**	**99.07**	*89.51*	*88.46*
Edge and sharpness	94.41	89.96	99.68	98.09
Frequency	99.67	98.36	99.31	97.42
Spatial organization	95.20	90.91	99.20	97.14
Residual and noise	99.33	97.60	96.60	93.46

Table 4: **Source classification with handcrafted features.** Accuracy (%) of an MLP trained on the full feature bank (top row) and on each category alone (Single), and accuracy of the full-bank MLP after permuting one category across test images (Perm.), on the held-out ImageNet-derived (ID) and Winoground-derived (OOD) sets. Among individual feature categories, texture achieves the highest single-category accuracy (bold) and the largest accuracy decrease under grouped permutation (italic).

Low-Level Visual Cues

Low-level visual cues distinguish T2I models. We construct a bank of 225 handcrafted features in seven categories, computed from each native 1024×1024 image: (i) **tone** (luminance statistics, dynamic range, entropy, and local contrast); (ii) **color** (RGB and perceptual chromaticity statistics, color histograms, inter-channel correlations, and colorfulness); (iii) **texture** (co-occurrence statistics, local binary patterns, and local entropy); (iv) **edge and sharpness** (gradient-orientation distributions, edge density, and high-frequency sharpness); (v) **frequency** (spectral energy distributions and oriented frequency responses); (vi) **spatial organization** (regional luminance and edge distributions, symmetry, and luminance centroids); and (vii) **residual and noise** (high-pass residual statistics, spatial autocorrelation, and residual spectral energy). The composition and definitions are given in Secs. C.1 and C.2. A lightweight MLP trained on the full bank (Sec. C.3) reaches 99.84% accuracy on the ID test set and 98.55% on the OOD set (Tab. 4). Low-level statistics alone thus support reliable source attribution across prompt distributions, with OOD accuracy close to ResNet-50's 98.94% (Fig. 4)

Texture features capture persistent model differences. We assess each category's standalone predictive power and grouped permutation importance with two experiments (Sec. C.4): training a separate MLP on one category alone, and permuting one category across test images while keeping the trained full-bank MLP and the other features fixed. Texture stands out on both measures (Tab. 4). Its 38 features alone reach 99.76% and 99.07% on the ID and OOD sets, respectively. Permuting texture reduces full-bank accuracy to 89.51% and 88.46%, producing the largest decrease among all feature categories, about ten percentage points on both sets. Frequency and residual features are also individually informative (99.67% and 99.33% on the ID set), but their permutation causes smaller accuracy decreases, at most

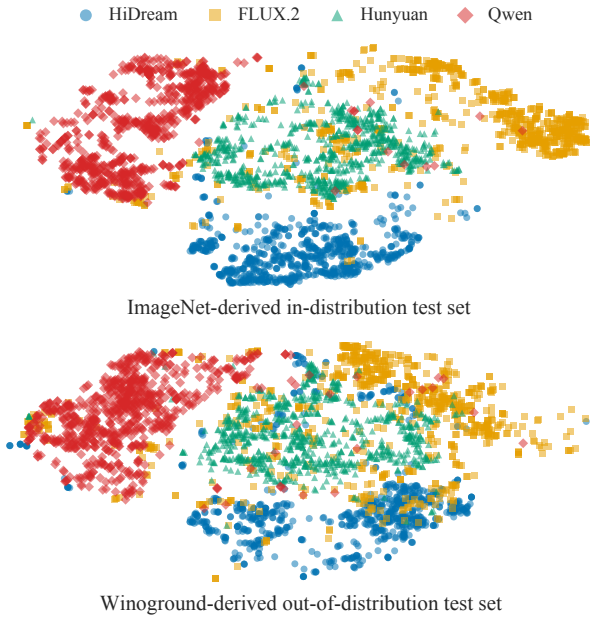

Figure 5: **t-SNE embedding of the 38 texture features.** A joint embedding of 3,200 images from each test set, shown separately for the ImageNet-derived and Winoground-derived sets. The model-dependent structure persists across the two prompt distributions.

5.1 percentage points. The texture category comprises eight Gray-Level Co-occurrence Matrix (GLCM) statistics (Haralick, Shanmugam, and Dinstein 1973), two uniform Local Binary Pattern (LBP) histograms (Ojala, Pietikainen, and Maenpaa 2002) with $(P, R) = (8, 1)$ and $(16, 2)$, and local-entropy statistics. A joint t-SNE embedding (Van der Maaten and Hinton 2008) of these features for 3,200 images from each test set shows clear model-dependent structure that persists across the two prompt distributions (Fig. 5). Texture is thus a prominent source of visual idiosyncrasies; visualizations of its four components and of the other six categories are in Secs. C.5 and C.6.

Mid-Level Visual Representations

Mid-level representations show model differences. We test whether the differences survive when an image is reduced to a mid-level representation. We estimate three: (i) **object layout** as segmentation masks from SAM 3.1 (Carion et al. 2026); (ii) **geometry** as depth maps from MoGe-2 (Wang et al. 2025a); and (iii) **appearance** as albedo maps from DiffusionRenderer (Liang et al. 2025a) (Sec. D.1; examples in Fig. A12). We train and evaluate a separate ResNet-50 on each representation. All three remain above chance on the ID set and lose only moderately on the OOD set (Tab. 5): even binary object masks reach 59.50%, depth reaches 87.72%, and albedo 99.75% (confusion matrices in Sec. D.3). The same ordering holds for the other classifier architectures (Sec. D.4).

Decoding model-specific mid-level visual structure. The representations also reveal what differs. For layout, we convert each mask into a signed distance map and average the

Representation	Acc. (ID)	Acc. (OOD)
Segmentation mask	59.50	46.73
Depth map	87.72	70.50
Albedo map	99.75	97.73

Table 5: **Source classification from mid-level representations.** Accuracy (%) of a separate ResNet-50 trained on each estimated representation. All three remain above the 25% chance level on both test sets, and albedo is nearly as informative as the original image.

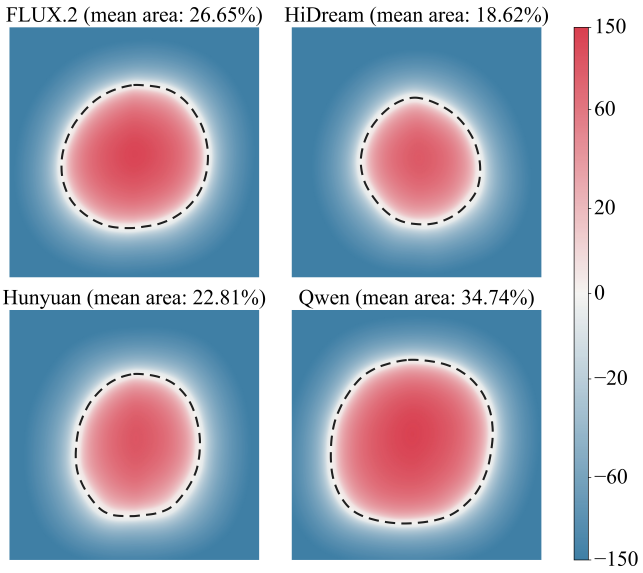

Figure 6: **Average object layouts.** Signed distance maps (in pixels) derived from segmentation masks on the ImageNet-derived set, averaged per model; dashed lines mark the zero-level contours. Qwen-Image has the largest average object.

maps per model over matched images (Sec. D.2); the averaged layouts differ visibly, and Qwen-Image in particular renders a substantially larger average object than the other models (Fig. 6). For geometry, the models occupy distinct ranges of object depth, summarized by the 5th and 95th depth percentiles of each image (Fig. 7). Depth interacts with the camera, so we also compare the focal lengths estimated by MoGe-2: HunyuanImage has markedly shorter focal lengths, *i.e.*, a wider field of view, which explains how its objects can be both nearer in depth and smaller in the image. For appearance, we measure the albedo map's fine-scale variation ratio, the share of albedo variation at fine rather than coarse spatial scales (Sec. D.2). Qwen-Image has the highest median fine-scale variation ratio (0.253), while HiDream and FLUX.2 have similarly lower medians (0.190 and 0.188), indicating more spatially coherent albedo regions. HunyuanImage is intermediate, with a median of 0.211 (Fig. 8). Mid-level representations thus offer a second, interpretable window on the visual idiosyncrasies of T2I models.

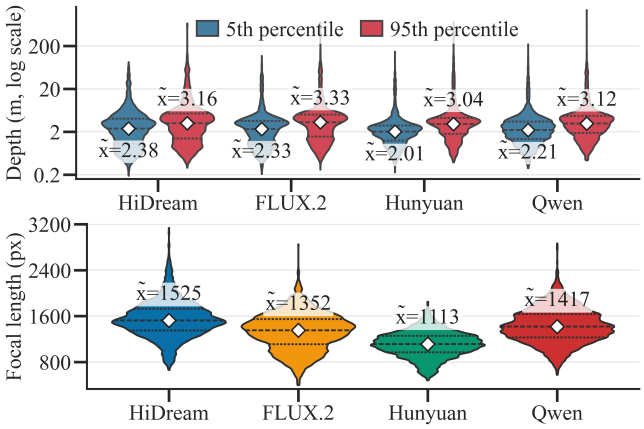

Figure 7: **Geometry differences.** Per-image 5th and 95th percentiles of object depth and the focal lengths estimated by MoGe-2. HunyuanImage combines the nearest objects with the shortest focal lengths.

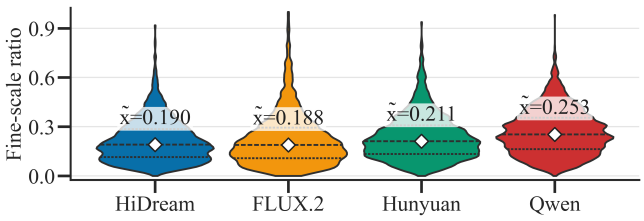

Figure 8: **Appearance differences.** Fine-scale variation ratio of the albedo map, the share of albedo variation that occurs at fine spatial scales; higher values indicate more fine-grained albedo changes.

High-Level Semantic Realization

High-level semantic realization reveals preferences of T2I models. Since all models receive identical prompts, any difference in how they realize the same concept reflects a preference of the model. We examine two aspects that the prompts leave unspecified, **visual style** and **background context**, by asking three Vision-Language Model (VLM) captioners, Qwen3-VL (Bai et al. 2025), InternVL3.5 (Wang et al. 2025b), and GLM-4.1V (Hong et al. 2025), to describe each image's style and background (Sec. E.1; example captions in Tab. A9). We represent each caption by word- and character-level TF-IDF features and train a logistic-regression classifier to predict the source model (Sec. E.2). Classification stays well above chance for both aspects and all three captioners, on the ID and OOD sets alike (Tab. 6; confusion matrices in Sec. E.3), so the models differ systematically in how they realize style and background.

Semantic realization patterns are model-specific. To see what differs, we fit a word-only TF-IDF classifier per captioner and list, for each model, the five unigrams or bigrams with the largest positive coefficients for background descriptions (Tab. 7; style in Sec. E.4). Several associations recur across all three captioners. HiDream backgrounds are associated with window and brick, FLUX.2 with warm,

Captioner	Prompt	Acc. (ID)	Acc. (OOD)
Qwen3-VL	Style	65.62	60.16
InternVL3.5		66.06	54.78
GLM-4.1V		68.78	58.19
Qwen3-VL	Background	71.12	65.11
InternVL3.5		70.94	60.53
GLM-4.1V		71.78	62.47

Table 6: **Source classification from VLM-generated captions.** Accuracy (%) of logistic-regression classifiers on TF-IDF features of style and background captions from three captioners; chance is 25%. Classification stays above chance for every captioner and both prompt distributions.

HiDream	FLUX.2	Hunyuan-Image	Qwen-Image
window	**warm**	hazy	**dark**
natural light	showing **warm**	**sunlit**	green
cloudy sky	golden	blurred with	is **dark**
brick	beige	bokeh	muted
partly cloudy	pale	weathered	green foliage
window	**warm**	sunlight	green
people	**warm** neutral	dusty	**dark**
brick	paved area	**sunlit**	doorway
walking	golden	branches suggesting	featuring green
focus setting	features light	soft	is **dark**
window	**warm**	**sunlit**	**dark**
window with	with **warm**	atmosphere	green
brick	smooth	hazy	is **dark**
flowers	neutral	soft	**dark** and
wooden	grassy area	muted	muted

Table 7: **Model-specific background preferences revealed by captions.** The five unigrams or bigrams with the largest positive coefficients of a word-level TF-IDF logistic-regression classifier for background descriptions; the three row blocks correspond, from top to bottom, to Qwen3-VL, InternVL3.5, and GLM-4.1V. Bold words rank among the top terms for the same model under all three captioners.

HunyuanImage with sunlit, and Qwen-Image with dark and green. These recurring terms characterize source-associated background realization under shared prompts. Individual coefficients describe predictive associations within the four-source roster between terms and sources.

Applications and Implications

We illustrate three applications and one implication of the visual idiosyncrasies beyond explaining the classifier.

A probe for guiding human judgment. We test whether the identified idiosyncrasies help people tell models apart. We repeat the human study of Fig. 2 with 61 new participants from the same pool and change only one thing: every question adds the hint "Please focus on comparing the images in terms of style, texture, and background," which summarizes the cues found at the low and high levels. The hint raises mean accuracy from 62.61% to 71.64% (Fig. 9); the 9.03-point gain

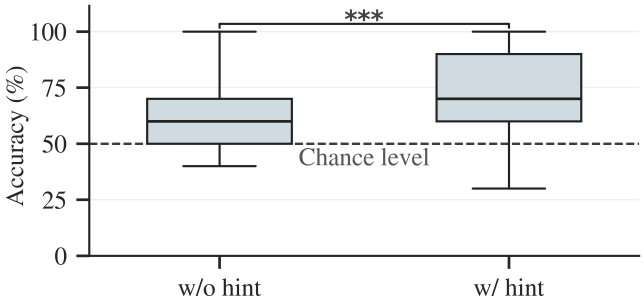

Figure 9: **Effect of a hint on human source classification.** Asking participants to compare style, texture, and background raises mean accuracy from 62.61% ($n = 69$) to 71.64% ($n = 61$). Statistical significance: $^*p < 0.001$ (stratified permutation test).

Held-out model	HiDream	FLUX.2	Hunyuan-Image	Qwen-Image
HiDream	-	**93.34**	3.02	<u>3.64</u>
FLUX.2	**58.29**	-	<u>31.28</u>	10.43
HunyuanImage	<u>48.49</u>	**49.37**	-	2.14
Qwen-Image	**76.88**	<u>20.23</u>	2.89	-

Table 8: **Inter-model similarity from leave-one-model-out classification.** Each row gives the distribution (%) of sources predicted for the held-out model's images by a ResNet-50 trained on the other three models. Bold marks the most frequent prediction and underline the second.

is significant under a stratified permutation test ($p < 0.001$) and positive in five of the six model pairs (Sec. B.5). The idiosyncrasies thus serve as interpretable probes that make model differences visible to people.

Quantifying inter-model visual similarity. We assess visual similarity by comparing sources through leave-one-source-out prediction. We exclude one source, train ResNet-50 on the other three, and report its predictions for the excluded source (Tab. 8). Held-out HiDream is predominantly assigned to FLUX.2 (93.34%). Held-out FLUX.2 is assigned to HiDream (58.29%) and HunyuanImage (31.28%). HunyuanImage divides almost equally between FLUX.2 (49.37%) and HiDream (48.49%), while Qwen-Image is most often assigned to HiDream (76.88%). These projections describe how each unseen source is partitioned by a classifier trained on the remaining sources.

Assessing real-synthetic similarity. We evaluate the trained source classifiers on 10,000 real ImageNet images, using the corresponding input representation for each classifier (Sec. A.4). Tab. 9 reports mean predicted probabilities. The RGB classifier predominantly assigns probability to FLUX.2 (89.16%), while the handcrafted-feature classifier favors HiDream (47.09%) followed by FLUX.2 (42.12%). Mask predictions are more dispersed, with HunyuanImage receiving 30.48%. Depth, albedo, style captions, and background captions favor FLUX.2. The responses differ across

Representation	HiDream	FLUX.2	Hunyuan-Image	Qwen-Image
Original	<u>7.37</u>	**89.16**	1.90	1.56
Handcrafted feature	**47.09**	<u>42.12</u>	1.41	9.38
Segmentation mask	21.66	21.48	**30.48**	<u>26.38</u>
Depth map	<u>26.53</u>	**36.87**	14.15	22.45
Albedo map	<u>26.68</u>	**54.69**	10.12	8.51
Caption (style)	20.88	**36.11**	13.63	<u>29.37</u>
Caption (background)	13.22	**39.55**	16.19	<u>31.04</u>

Table 9: **Mean predicted source probabilities (%) for real ImageNet images across representation levels.** Each row averages the probabilities produced by the source classifier trained on the corresponding representation. Bold marks the largest probability and underline the second largest.

representation levels and characterize the behavior of these source-trained classifiers on real images.

Risks to anonymous evaluation. Persistent visual idiosyncrasies challenge the anonymity assumed by human preference evaluations such as the Text-to-Image Arena (Arena 2026). Even when model identities are hidden, the images themselves may reveal the source, and a participant can use a trained source classifier to infer a generator and vote for or against that identity. Our human study addresses a separate question: whether a hint about the characterized cues improves human source identification; it is not an empirical test of an arena attack. Naseh et al. (2026) demonstrate leaderboard deanonymization through prompt-conditioned image-embedding clusters; our results complement this attack-oriented view by characterizing, for more recent models, the multi-level cues that make deanonymization possible.

Conclusion

Images from four frontier T2I models remain almost perfectly distinguishable by source under identical prompts, across classifier architectures, under severe corruptions, and on out-of-distribution prompts, even though humans distinguish them only modestly. This robustness suggests that the distinguishing signal reflects stable properties of each model's learned output distribution rather than superficial artifacts of prompting, resolution, or post-processing. We trace this distinguishability to visual idiosyncrasies at three levels: texture statistics that alone suffice for classification, systematic differences in object layout, depth, focal length, and albedo, and model-specific style and background preferences that captions expose. Together, these findings indicate that frontier models with similar benchmark scores still occupy noticeably different regions of image space. These idiosyncrasies make source classification an interpretable probe of how generative models differ, guide human judgment, quantify inter-model and real-synthetic similarity, and warn against assuming anonymity in model evaluation. Limitations and future directions are discussed in Sec. G.

Acknowledgments

This work is supported in part by the: Brain Science and Brain-like Intelligence Technology-National Science and Technology Major Project (2026ZD0222200), the National Natural Science Foundation of China (62376009), the Beijing Nova program, the NVIDIA Academic Grant Program using Spark and Thor, the PKU-BingJi Joint Laboratory for Artificial Intelligence, the Wuhan Major Scientific and Technological Special Program (2025060902020304), the Hubei Embodied Intelligence Foundation Model Research and Development Program, and the National Comprehensive Experimental Base for Governance of Intelligent Society, Wuhan East Lake High-Tech Development Zone.

References

Amoroso, R.; Morelli, D.; Cornia, M.; Baraldi, L.; Del Bimbo, A.; and Cucchiara, R. 2024. Parents and children: Distinguishing multimodal deepfakes from natural images. *ACM Transactions on Multimedia Computing, Communications and Applications (TOMM)*, 21(1): 1–23.

Arena. 2026. Text-to-Image Arena. https://arena.ai/leaderboard/text-to-image.

Bai, S.; Cai, Y.; Chen, R.; Chen, K.; Chen, X.; Cheng, Z.; Deng, L.; Ding, W.; Gao, C.; Ge, C.; et al. 2025. Qwen3-vl technical report. *arXiv preprint arXiv:2511.21631*.

Black Forest Labs. 2025. FLUX.2: Frontier Visual Intelligence. https://bfl.ai/blog/flux-2.

Bradski, G.; and Kaehler, A. 2008. *Learning OpenCV: Computer vision with the OpenCV library*. " O'Reilly Media, Inc.".

Breiman, L. 2001. Random forests. *Machine Learning*, 45(1): 5–32.

Brokman, J.; Giloni, A.; Hofman, O.; Vainshtein, R.; Kojima, H.; and Gilboa, G. 2025. Manifold Induced Biases for Zero-Shot and Few-Shot Detection of Generated Images. In *International Conference on Learning Representations (ICLR)*.

Bui, T.; Yu, N.; and Collomosse, J. 2022. RepMix: Representation Mixing for Robust Attribution of Synthesized Images. In *European Conference on Computer Vision (ECCV)*.

Cai, Q.; Chen, J.; Gao, C.; Gong, Z.; Li, Y.; Pan, Y.; Peng, Y.; Qiu, Z.; Yu, K.; Zhang, Y.; et al. 2026. Hidream-o1-image: A natively unified image generative foundation model with pixel-level unified transformer. *arXiv preprint arXiv:2605.11061*.

Canny, J. 1986. A Computational Approach to Edge Detection. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 8(6): 679–698.

Cao, S.; Chen, H.; Chen, P.; Cheng, Y.; Cui, Y.; Deng, X.; Dong, Y.; Gong, K.; Gu, T.; Gu, X.; et al. 2025. HunyuanImage 3.0 Technical Report. *arXiv preprint arXiv:2509.23951*.

Carion, N.; Gustafson, L.; Hu, Y.-T.; Debnath, S.; Hu, R.; Suris Coll-Vinent, D.; Ryali, C.; Alwala, K. V.; Khedr, H.; Huang, A.; et al. 2026. SAM 3: Segment Anything with Concepts. In *International Conference on Learning Representations (ICLR)*.

Chai, L.; Bau, D.; Lim, S.-N.; and Isola, P. 2020. What Makes Fake Images Detectable? Understanding Properties That Generalize. In *European Conference on Computer Vision (ECCV)*.

Choi, S.; Lee, H.; Lee, J.; Kim, R.; Choi, S. J.; and Lee, M. 2026. A Debiased Reconstruction-Based Framework for Training-Free Detection of AI-Generated Images. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Choi, S.; Lee, H.; and Lee, M. 2025. Training-Free Detection of AI-Generated Images via Cropping Robustness. In *Advances in Neural Information Processing Systems (NeurIPS)*.

Corvi, R.; Cozzolino, D.; Poggi, G.; Nagano, K.; and Verdoliva, L. 2023a. Intriguing properties of synthetic images: from generative adversarial networks to diffusion models. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Corvi, R.; Cozzolino, D.; Zingarini, G.; Poggi, G.; Nagano, K.; and Verdoliva, L. 2023b. On the detection of synthetic images generated by diffusion models. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.

Cozzolino, D.; Poggi, G.; Corvi, R.; Nießner, M.; and Verdoliva, L. 2024a. Raising the bar of ai-generated image detection with clip. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Cozzolino, D.; Poggi, G.; Nießner, M.; and Verdoliva, L. 2024b. Zero-Shot Detection of AI-Generated Images. In *European Conference on Computer Vision (ECCV)*.

Dalal, N.; and Triggs, B. 2005. Histograms of oriented gradients for human detection. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Frank, J.; Eisenhofer, T.; Schönherr, L.; Fischer, A.; Kolossa, D.; and Holz, T. 2020. Leveraging frequency analysis for deep fake image recognition. In *International Conference on Machine Learning (ICML)*.

Gabor, D. 1946. Theory of communication. Part 1: The analysis of information. *Institution of Electrical Engineers-part III: Radio and Communication engineering*, 93(26): 429–441.

Girish, S.; Suri, S.; Rambhatla, S. S.; and Shrivastava, A. 2021. Towards Discovery and Attribution of Open-World GAN Generated Images. In *International Conference on Computer Vision (ICCV)*.

Gregorutti, B.; Michel, B.; and Saint-Pierre, P. 2015. Grouped variable importance with random forests and application to multiple functional data analysis. *Computational Statistics & Data Analysis*, 90: 15–35.

Guarnera, L.; Giudice, O.; and Battiato, S. 2024. Level up the deepfake detection: a method to effectively discriminate images generated by gan architectures and diffusion models. In *Intelligent Systems Conference (IntelliSys)*.

Haralick, R. M.; Shanmugam, K.; and Dinstein, I. H. 1973. Textural features for image classification. *IEEE Transactions on Systems, Man, and Cybernetics*, 3(6): 610–621.

Harris, C. R.; Millman, K. J.; Van Der Walt, S. J.; Gommers, R.; Virtanen, P.; Cournapeau, D.; Wieser, E.; Taylor, J.; Berg, S.; Smith, N. J.; et al. 2020. Array programming with numpy. *Nature*, 585(7825): 357–362.

Hasler, D.; and Suesstrunk, S. E. 2003. Measuring colorfulness in natural images. In *Human vision and electronic imaging VIII*.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Hong, W.; Yu, W.; Gu, X.; Wang, G.; Gan, G.; Tang, H.; Cheng, J.; Qi, J.; Ji, J.; Pan, L.; et al. 2025. Glm-4.5v and glm-4.1 v-thinking: Towards versatile multimodal reasoning with scalable reinforcement learning. *arXiv preprint arXiv:2507.01006*.

Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations (ICLR)*.

Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*.

Laszkiewicz, M.; Ricker, J.; Lederer, J.; and Fischer, A. 2024. Single-Model Attribution of Generative Models Through Final-Layer Inversion. In *International Conference on Machine Learning (ICML)*.

Lee, J.; Koo, M.; and Kwak, N. 2026. CSF: Black-Box Fingerprinting via Compositional Semantics for Text-to-Image Models. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Liang, R.; Gojcic, Z.; Ling, H.; Munkberg, J.; Hasselgren, J.; Lin, C.-H.; Gao, J.; Keller, A.; Vijaykumar, N.; Fidler, S.; et al. 2025a. Diffusion renderer: Neural inverse and forward rendering with video diffusion models. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Liang, Y.; Yu, M.; Li, G.; Jiang, J.; Du, F.; Li, J.; Xie, L.; Xu, Z.; and Huang, W. 2025b. Denoising Trajectory Biases for Zero-Shot AI-Generated Image Detection. In *Advances in Neural Information Processing Systems (NeurIPS)*.

Liu, D. C.; and Nocedal, J. 1989. On the limited memory BFGS method for large scale optimization. *Mathematical Programming*, 45(1): 503–528.

Liu, F.; Luo, H.; Li, Y.; Torr, P.; and Gu, J. 2024a. Which model generated this image? a model-agnostic approach for origin attribution. In *European Conference on Computer Vision (ECCV)*.

Liu, H.; Tan, Z.; Tan, C.; Wei, Y.; Wang, J.; and Zhao, Y. 2024b. Forgery-Aware Adaptive Transformer for Generalizable Synthetic Image Detection. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Liu, Z.; and He, K. 2025. A decade's battle on dataset bias: Are we there yet? In *International Conference on Learning Representations (ICLR)*.

Liu, Z.; Qi, X.; and Torr, P. H. 2020. Global texture enhancement for fake face detection in the wild. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Luo, M. R. 2023. Cielab. In *Encyclopedia of Color Science and Technology*, 251–257. Springer.

Marra, F.; Gragnaniello, D.; Verdoliva, L.; and Poggi, G. 2019. Do gans leave artificial fingerprints? In *Conference on Multimedia Information Processing and Retrieval (MIPR)*.

Naseh, A.; Suri, A.; Peng, Y.; Chaudhari, H.; Oprea, A.; and Houmansadr, A. 2026. When Anonymity Breaks: Identifying Models Behind Text-to-Image Leaderboards. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Nguyen, T. D.; Azizpour, A.; and Stamm, M. C. 2025. Forensic Self-Descriptions Are All You Need for Zero-Shot Detection, Open-Set Source Attribution, and Clustering of AI-Generated Images. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Ojala, T.; Pietikainen, M.; and Maenpaa, T. 2002. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 24(7): 971–987.

Ojha, U.; Li, Y.; and Lee, Y. J. 2023. Towards universal fake image detectors that generalize across generative models. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2024. DINOv2: Learning Robust Visual Features without Supervision. *Transactions on Machine Learning Research (TMLR)*.

Otsu, N. 1979. A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*, 9(1): 62–66.

Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems (NeurIPS)*.

Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research (JMLR)*, 12: 2825–2830.

Qwen Team. 2025. Qwen-Image-2512 Model Card. https://huggingface.co/Qwen/Qwen-Image-2512.

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. 2021. Learning Transferable Visual Models from Natural Language Supervision. In *International Conference on Machine Learning (ICML)*.

Ricker, J.; Lukovnikov, D.; and Fischer, A. 2024. AEROBLADE: Training-Free Detection of Latent Diffusion Images Using Autoencoder Reconstruction Error. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.;

Berg, A. C.; and Fei-Fei, L. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3): 211–252.

Salton, G.; and Buckley, C. 1988. Term-weighting approaches in automatic text retrieval. *Information Processing & Management*, 24(5): 513–523.

Sarkar, A.; Mai, H.; Mahapatra, A.; Lazebnik, S.; Forsyth, D. A.; and Bhattad, A. 2024. Shadows Don't Lie and Lines Can't Bend! Generative Models Don't Know Projective Geometry...for Now. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Sha, Z.; Li, Z.; Yu, N.; and Zhang, Y. 2023. De-fake: Detection and attribution of fake images generated by text-to-image generation models. In *ACM Conference on Computer and Communications Security (CCS)*.

Shannon, C. E. 1948. A mathematical theory of communication. *The Bell System Technical Journal*, 27(3): 379–423.

Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research (JMLR)*, 15(1): 1929–1958.

Sun, M.; Yin, Y.; Xu, Z.; Kolter, J. Z.; and Liu, Z. 2025. Idiosyncrasies in Large Language Models. In *International Conference on Machine Learning (ICML)*.

Thrush, T.; Jiang, R.; Bartolo, M.; Singh, A.; Williams, A.; Kiela, D.; and Ross, C. 2022. Winoground: Probing Vision and Language Models for Visio-Linguistic Compositionality. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Torralba, A.; and Efros, A. A. 2011. Unbiased look at dataset bias. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of Machine Learning Research (JMLR)*, 9(86): 2579–2605.

Van der Walt, S.; Schönberger, J. L.; Nunez-Iglesias, J.; Boulogne, F.; Warner, J. D.; Yager, N.; Gouillart, E.; and Yu, T. 2014. scikit-image: image processing in Python. *PeerJ*, 2: e453.

Wang, C.; Yang, Z.; Wang, Y.; Zhang, W.; and Chen, K. 2026. AEDR: Training-Free AI-Generated Image Attribution via Autoencoder Double-Reconstruction. In *AAAI Conference on Artificial Intelligence (AAAI)*.

Wang, R.; Xu, S.; Dong, Y.; Deng, Y.; Xiang, J.; Lv, Z.; Sun, G.; Tong, X.; and Yang, J. 2025a. Moge-2: Accurate monocular geometry with metric scale and sharp details. In *Advances in Neural Information Processing Systems (NeurIPS)*.

Wang, S.-Y.; Wang, O.; Zhang, R.; Owens, A.; and Efros, A. A. 2020. CNN-Generated Images Are Surprisingly Easy to Spot... for Now. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Wang, W.; Gao, Z.; Gu, L.; Pu, H.; Cui, L.; Wei, X.; Liu, Z.; Jing, L.; Ye, S.; Shao, J.; et al. 2025b. Internvl3. 5: Advancing open-source multimodal models in versatility, reasoning, and efficiency. *arXiv preprint arXiv:2508.18265*.

Wang, Z.; Bao, J.; Zhou, W.; Wang, W.; Hu, H.; Chen, H.; and Li, H. 2023a. Dire for diffusion-generated image detection. In *International Conference on Computer Vision (ICCV)*.

Wang, Z.; Chen, C.; Zeng, Y.; Lyu, L.; and Ma, S. 2023b. Where did i come from? origin attribution of ai-generated images. In *Advances in Neural Information Processing Systems (NeurIPS)*.

Wang, Z.; Sehwag, V.; Chen, C.; Lyu, L.; Metaxas, D. N.; and Ma, S. 2024. How to Trace Latent Generative Model Generated Images without Artificial Watermark? In *International Conference on Machine Learning (ICML)*.

Xu, K.; Zhang, L.; and Shi, J. 2025. Detecting Origin Attribution for Text-to-Image Diffusion Models. In *Proceedings of Winter Conference on Applications of Computer Vision (WACV)*.

Yu, N.; Davis, L. S.; and Fritz, M. 2019. Attributing Fake Images to GANs: Learning and Analyzing GAN Fingerprints. In *International Conference on Computer Vision (ICCV)*.

Zeng, B.; Yin, Y.; and Liu, Z. 2024. Understanding bias in large-scale visual datasets. In *Advances in Neural Information Processing Systems (NeurIPS)*.

A Experimental Details

A.1 Prompt Construction

All source models receive the same prompts, so that prompt content is controlled. We use two prompt families: ImageNet-derived prompts, which name a single object category, and Winoground-derived prompts, which describe compositional relations among multiple entities. Example images for both families are shown in Fig. A10.

ImageNet-derived prompts. We construct one prompt for each of the 1,000 ImageNet-1K classes (Russakovsky et al. 2015). For a class label with a primary name followed by comma-separated aliases, the primary name follows "a photo of" and the aliases are appended after "which is also known as"; the second clause is omitted when a class has no alias. Four seeds per prompt yield 4,000 images per source model.

ImageNet-derived prompt template and examples

Template: `a photo of {primary class name} [, which is also known as {alias 1}, {alias 2}, ...]`

Examples:

```
a photo of stingray
a photo of tabby, which is also known as tabby cat
a photo of goldfish, which is also known as Carassius auratus
a photo of toilet tissue, which is also known as toilet paper, bathroom tissue
```

Winoground-derived prompts. Winoground (Thrush et al. 2022) contains 400 examples, each a contrastive pair (`caption_0`, `caption_1`) that reverses or alters a compositional relation. We use `caption_0` verbatim as the generation prompt, without any photographic or stylistic wrapper; `caption_1` is not used. The same four seeds yield 1,600 images per source model, which form the out-of-distribution (OOD) test set.

Winoground-derived prompt template and examples

Template: `{caption_0}`

Examples:

```
an old person kisses a young person
a bird eats a snake
a bottle is in water
light colored dog with dark colored cat
```

A.2 Source Classifiers and Training

Architectures. ResNet-50 (He et al. 2016) is trained from random initialization with its final fully connected layer replaced by an N-way classification head. The two-layer Convolutional Neural Network (CNN) uses two 3×3 convolutional layers with 32 and 64 output channels, each followed by ReLU and 2×2 max pooling; the resulting feature map is adaptively average-pooled to 8×8, flattened, and passed through a two-layer head with dimensions $4096 \to 128 \to N$ and an intermediate ReLU. For DINOv2 ViT-B/14 (Oquab et al. 2024) and CLIP ViT-B/32 (Radford et al. 2021), we freeze the visual encoders and train a single linear classifier on their image features.

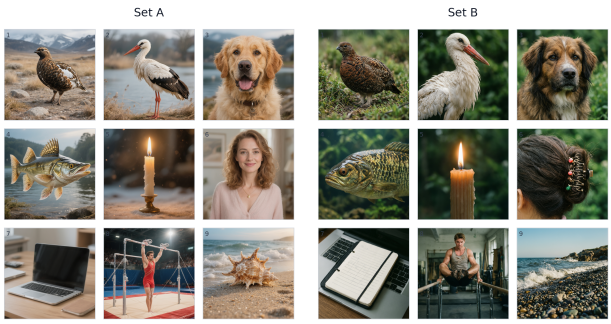

Figure A1: **Example reference board used in the human study.** Each anonymous set contains nine examples from one source model, with matched grid positions generated from the same prompt. The assignment of models to Set A and Set B is counterbalanced across study versions.

Data splits. We partition the 1,000 ImageNet-derived classes into 800 training classes (200–999) and 200 in-distribution (ID) evaluation classes (0–199). With four generation seeds and four source models, this yields 12,800 training and 3,200 ID evaluation images, balanced across sources. No separate validation split is used: checkpoints are selected by accuracy on the training set. Feature normalization is fitted only on the training set. All 6,400 Winoground-derived images form the OOD test set and are used only for final evaluation, without affecting normalization or checkpoint selection.

Training. Each image is decoded, converted to RGB, and resized to 224×224 with bicubic interpolation. ResNet-50 and two-layer CNN inputs are normalized channel-wise with mean 0.5 and standard deviation 0.5; DINOv2 and CLIP inputs use their respective model-specific preprocessing. We apply no stochastic data augmentation. All classifiers are trained with cross-entropy loss and Adam (Kingma and Ba 2015) at a fixed learning rate of 10^{-4} for 100 epochs with a batch size of 16. ResNet-50 and the two-layer CNN are optimized end-to-end, whereas the DINOv2 and CLIP encoders remain frozen. We select the checkpoint with the highest training accuracy, and report its accuracy on the held-out test sets. All experiments run on a single NVIDIA H100 GPU.

A.3 Human Study Protocol

Participants and design. We recruit adult participants through Prolific and administer the study in English using Google Forms. Participants provide informed consent, confirm that they are at least 18 years old and hold at least a bachelor's degree, and are instructed to use a desktop or laptop computer.

The four source models form six unordered model pairs. Each participant is assigned to one pair and learns to distinguish its two sources, anonymously labeled Set A and Set B. For each pair, we counterbalance which model is assigned to which set and cross this mapping with two instruction conditions, with and without a hint, yielding $6 \times 2 \times 2 = 24$ between-participant study versions. Repeat participation is



(a) ResNet-50.

(b) Two-layer CNN.

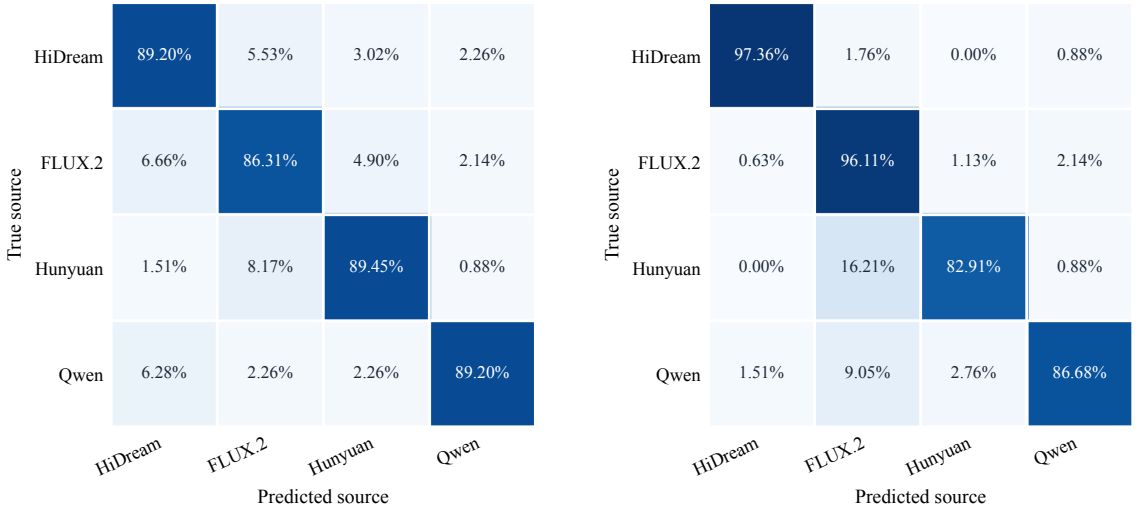

(c) Linear classifier on DINOv2 features.

(d) Linear classifier on CLIP features.

Figure A2: **Source-classification confusion matrices on the held-out ImageNet-derived test set.** Rows denote true source models and columns denote predicted sources. Each cell reports the row-normalized percentage.

prevented through Prolific. New participants are automatically assigned to one of the currently least-filled versions, and returning participants are directed back to their original version.

We receive 150 submissions and exclude participants who answer either of two attention checks incorrectly, leaving 130 participants: 69 in the no-hint condition, analyzed in the main text, and 61 in the hint condition. All retained participants complete the ten main trials and both attention checks.

Stimuli. All stimuli are sampled reproducibly from the ImageNet-derived set. The same nine reference prompts and ten main-trial prompts are used across all model pairs, with no overlap between reference and trial prompts. Within each pair, the two models are always compared on images generated from matched prompts. Each reference board shows nine examples from each source as two 3×3 grids, with matched grid positions generated from the same prompt (Fig. A1). The two mapping versions use identical images and differ only in which source is labeled Set A. The ten main trials follow a fixed prompt order and contain five query images from each source.

Procedure. Participants complete one trial per page. Each page displays the assigned reference board, a query image with its prompt, and the question *Which reference set is this image more likely to come from?* In the hint condition, the sentence *Hint: Please focus on comparing the images in terms of style, texture, and background.* is added to every question. Participants select Set A or Set B; model identities and correct answers are never shown. The sequence consists of three main trials, one attention check, three main trials, one

Setting (s, g)	FLUX.2-dev (FLUX.2) recall	Overall acc.
$(50, 4)$	99.75	99.94
$(25, 4)$	99.87	99.97
$(75, 4)$	99.87	99.97
$(50, 2)$	99.75	99.94
$(50, 6)$	99.87	99.97

Table A1: **Source attribution across FLUX.2 generation settings.** Here s is the number of inference steps and g is distilled guidance. A single frozen four-way ResNet-50 evaluates 800 FLUX.2 images per setting. Overall accuracy also includes 800 matched images from each of HiDream-O1-Dev (HiDream), HunyuanImage-3.0 (HunyuanImage), and Qwen-Image-2512 (Qwen-Image).

Comparison	Centroid	MMD^2
FLUX.2 $(25, 4)$	0.0148	0.000167
FLUX.2 $(75, 4)$	0.0076	0.000045
FLUX.2 $(50, 2)$	0.0248	0.000469
FLUX.2 $(50, 6)$	0.0136	0.000144
HiDream	0.1688	0.021128
HunyuanImage	0.1587	0.018803
Qwen-Image	0.1722	0.022227

Table A2: **CLIP distances within FLUX.2 and across generators.** Every row compares the indicated group with the newly generated FLUX.2 $(50, 4)$ reference. The four FLUX.2 rows vary generation settings, and the final three rows compare different generators. Each group contains 800 images matched by prompt and generation seed. Smaller values indicate closer groups in CLIP feature space.

attention check, and four main trials. Each attention check reuses an image from the reference board and is presented as an ordinary trial. The primary outcome is participant-level mean accuracy over the ten main trials, with chance at 50%; attention checks and a final cue-ranking question are excluded from this computation.

A.4 Real ImageNet Images

For the real-synthetic similarity analysis (Tab. 9), we sample 10,000 real ImageNet images covering all 1,000 classes, with one to eleven images per class. Each real image passes through the same pipelines as the generated images: the handcrafted feature extractor of Sec. C, SAM 3.1 queried with the class label, MoGe-2, DiffusionRenderer, and Qwen3-VL-8B-Instruct as the Vision-Language Model (VLM) captioner under both prompts of Sec. E.1. Each representation is then fed to the four-way classifier trained on that representation. For the segmentation-mask representation, the mask-area filter of Sec. D.2 retains 7,653 images from 982 classes. The sampled real images are originally square. We convert them to RGB and resize them to 1024×1024 using Lanczos interpolation, saving them as JPEG with quality 95 and no chroma subsampling. Handcrafted features are extracted directly from these prepared images without further resizing or cropping.

Corruption	Setting	Acc.
None	-	100.00
Color jittering	Strength: 1.0	100.00
Color jittering	Strength: 2.0	99.94
Gaussian blur	Radius: 1.0	100.00
Gaussian blur	Radius: 2.0	100.00
Gaussian blur	Radius: 4.0	100.00
Gaussian blur	Radius: 8.0	100.00
Gaussian noise	Blending weight: 0.2	100.00
Gaussian noise	Blending weight: 0.4	99.91
Gaussian noise	Blending weight: 0.6	99.59
Gaussian noise	Blending weight: 0.8	97.99
Low resolution	Resolution: 64×64	100.00
Low resolution	Resolution: 32×32	100.00
Low resolution	Resolution: 16×16	99.94
Low resolution	Resolution: 8×8	99.21
Low resolution	Resolution: 4×4	95.60
Low resolution	Resolution: 2×2	84.74
Patch shuffle	Patch size: 32×32	99.43
Patch shuffle	Patch size: 16×16	97.02
Patch shuffle	Patch size: 8×8	92.49
Reconstruction	SD-VAE	100.00
Reconstruction	Qwen-VAE	100.00

Table A3: **Full corruption and transformation results.** For each setting, a ResNet-50 is trained and evaluated on the same transformed ImageNet-derived distribution. Accuracy remains high under moderate corruptions, stays well above chance under the most severe settings, and is essentially unaffected by patch shuffling with large patches or by VAE reconstruction.

B Additional Results on Source Classification

B.1 Confusion Matrices

Figure A2 reports the row-normalized confusion matrices of the four classifiers on the held-out ImageNet-derived test set. The two end-to-end CNNs produce strongly diagonal matrices: ResNet-50 reaches 100.00% accuracy and the two-layer CNN 98.81%, with the latter's residual errors concentrated in FLUX.2 images misclassified as HiDream. Linear classifiers on frozen DINOv2 and CLIP features remain well above the 25% chance level but show broader confusion: DINOv2 most often confuses FLUX.2 with HiDream, whereas CLIP assigns images from several sources to FLUX.2. Source identity is thus most readily recovered by end-to-end visual classifiers, while frozen semantic features retain a weaker source signal.

B.2 Robustness to FLUX.2 Generation Settings

We evaluate FLUX.2 under five settings of inference steps and distilled guidance. Each setting uses the same 200 ImageNet-derived evaluation prompts and generation seeds 42, 66, 222, and 999, yielding 800 images at 1024×1024 resolution. Each FLUX.2 batch is evaluated alongside the other three generators using the frozen four-way ResNet-50 from the main experiments with HiDream. Table A1 shows

Task	Classifier	Acc.
Four-way	ResNet-50	69.58 ± 3.66
Four-way	Full MLP (225D)	97.92 ± 0.36
Four-way	Texture MLP (38D)	98.33 ± 0.36
Eight-way	Texture MLP (38D)	97.29 ± 0.18

Table A4: **Source attribution for API-generated images.** Four-way classification uses the four API services, and eight-way classification adds the four main open-source generators. Each source contributes 157 training images and 40 validation images from disjoint categories. Values are validation accuracies in percent, reported as means $\pm$ sample standard deviations across training seeds 42, 43, and 44 at the final epoch. The ResNet-50 is randomly initialized.

Model pair	n (no-hint/hint)	Effect
FLUX.2–HiDream	17/11	9.57
FLUX.2–HunyuanImage	10/9	12.11
FLUX.2–Qwen-Image	10/10	5.00
HiDream–HunyuanImage	11/10	10.00
HiDream–Qwen-Image	11/10	10.55
HunyuanImage–Qwen-Image	10/11	-0.82

Table A5: **Pair-specific hint effects.** Effects are hint minus no-hint mean accuracy in percentage points. Five of the six pair-specific effects are positive.

that FLUX.2 remains reliably identifiable across the tested settings.

We additionally compare image distributions using unit-normalized frozen CLIP ViT-B/32 features. Centroid distance is the Euclidean distance between group means. Gaussian-kernel MMD^2, with bandwidth $\sigma = 1$, also captures distributional differences beyond the mean. Using the newly generated FLUX.2 $(50, 4)$ batch as the reference, Tab. A2 shows smaller differences for the four setting variants than for the three other generators.

B.3 Full Corruption Sweeps and Additional Transformations

Table A3 reports all settings of the four corruptions in the main text, of which Tab. 3 shows a subset. Each transformation is applied to both the training and the held-out images, and a separate ResNet-50 is trained and evaluated on each transformed distribution. Gaussian blur is implemented with the `GaussianBlur` filter of Pillow, whose only parameter is the blur radius in pixels; the kernel size follows from the radius. Examples of corrupted images are shown in Fig. A11.

We additionally consider two transformations. Patch shuffle divides each image into non-overlapping $p \times p$ patches and randomly permutes their spatial locations; classifiers are trained and evaluated on shuffled images. Accuracy decreases from 99.43% with 32×32 patches to 92.49% with 8×8 patches. Local image statistics thus remain highly predictive after global spatial organization is disrupted, while increasingly local shuffling removes useful structure. Reconstruction encodes and decodes every image with the same off-the-shelf Variational Autoencoder (VAE) before training and evaluation. ResNet-50 reaches 100.00% with both the Stable Diffusion VAE (Rombach et al. 2022) and the Qwen-Image VAE (Qwen Team 2025). The source signal therefore survives a shared learned reconstruction bottleneck and cannot be explained solely by fragile final-stage decoder artifacts.

B.4 Source Attribution for Closed-Source Models

We extend the experiment to GPT-Image-2.5 Sunburst and Qwen-Image-3.0-Pro through their official APIs, and Gemini Nano Banana 2.1 and Grok Imagine Image 2.0 through Poe. We reuse the original class prompts and collect native 1024×1024 outputs, with GPT quality set to high, Grok quality set to medium, and Qwen prompt expansion and visible watermark disabled in the request. We retain 197 categories with successful outputs from all four services, using one image per source and category. A fixed random partition assigns 157 categories to training and 40 to validation, shared across sources. Classifiers are trained from scratch for 100 epochs, and feature normalization is fitted on training images. For the eight-way extension, we add one matched generation-seed-42 image per category from each of the four main open-source generators, including HiDream. Table A4 shows strong texture-based attribution in both the four-way and eight-way tasks, extending the evidence to native API outputs under these service configurations.

B.5 Statistical Analysis of the Human Study

Effect of the hint. To test the effect of the hint while respecting the blocked design, we treat each participant's mean accuracy over the ten main trials as one observation and use a two-sided stratified permutation test. Under the null hypothesis of no hint effect, we independently permute the hint and no-hint labels within each model-pair-by-mapping block, preserving the observed condition counts in every block. Across 50,000 permutations, the observed difference of 9.03 percentage points (71.64% with the hint versus 62.61% without) has $p < 0.001$. A stratified nonparametric bootstrap that resamples participants within each model-pair-by-mapping-by-condition cell (50,000 resamples) gives a 95% confidence interval of $[4.76, 13.25]$ percentage points.

As a trial-level robustness check, we fit a Bayesian mixed-effects logistic regression to the 1,300 binary main-trial responses, with hint condition, model pair, mapping, and centered trial index as fixed effects and random intercepts for participant and query image. With weakly informative $\mathcal{N}(0, 5^2)$ priors on the fixed effects and $\mathcal{N}(0, 1)$ priors on the log random-effect standard deviations, the estimated odds ratio of the hint is 1.51, with a 95% credible interval of $[1.25, 1.82]$.

Consistency across pairs and exclusion rules. Table A5 reports the hint effect within each model pair. Five of the six effects are positive, so the overall effect is not driven by a single pair. The result is also robust to the exclusion rule: excluding only the participant who fails both attention checks retains 149 participants (73 hint and 76 no-hint) and yields a difference of 8.45 percentage points (stratified permutation

Category	Dim.	Descriptor composition
Tone	32	Four channel-average intensity moments; 16-bin HSV-value histogram; seven luminance quantiles; two clipping ratios; luminance entropy; two local RMS-contrast statistics.
Color	60	Twelve RGB moments; 16-bin hue and saturation histograms; eight CIELAB a^*, b^* moments; three RGB correlations; colorfulness; mean sine and cosine of hue; two Lab-chroma statistics.
Texture	38	Eight GLCM statistics; 10-bin $(P, R) = (8, 1)$ and 18-bin $(P, R) = (16, 2)$ uniform-LBP histograms; two local-entropy statistics.
Edge and sharpness	20	Three Canny densities; three gradient-magnitude statistics; four HOG summaries; 9-bin magnitude-weighted Sobel-orientation histogram; Laplacian variance.
Frequency	21	FFT high/low-energy ratio; four radial energy bands; 12 mean Gabor amplitudes; spectral slope, centroid, spread, and flatness.
Spatial organization	40	Nine Otsu-mask morphology statistics; 18 cell-wise luminance statistics and nine edge densities on a 3×3 grid; two reflection errors; two luminance-centroid coordinates.
Residual and noise	14	Four high-pass residual moments; three absolute-residual quantiles; three one-pixel autocorrelations; four radial energy bands.
Total	**225**	**Seven mutually exclusive feature categories.**

Table A6: **Composition of the 225-dimensional handcrafted feature bank.** Features are extracted from each native 1024×1024 image. "Dim." denotes the number of scalar descriptors assigned to each mutually exclusive category. All descriptors are computed per image before standardization with ImageNet-derived training statistics.

$p < 0.001$), and including all 150 submissions (73 hint and 77 no-hint) yields 8.61 percentage points ($p < 0.001$).

C Low-Level Visual Cues

C.1 Feature Bank Composition

We extract a 225-dimensional handcrafted feature vector from every image directly from the original 1024×1024 PNG, without resizing or cropping. The bank comprises seven disjoint categories: tone (32 dimensions), color (60), texture (38), edge and sharpness (20), frequency (21), spatial organization (40), and residual and noise (14). A fixed registry assigns every descriptor to exactly one category (Tab. A6).

C.2 Feature Definitions

Shared definitions. For normalized color channels $R, G, B \in [0, 1]$, we define the channel-average intensity I and the luminance Y as

$$\begin{aligned} I &= (R + G + B)/3, \\ Y &= 0.2126R + 0.7152G + 0.0722B. \end{aligned} \tag{A1}$$

For values $x_1, \ldots, x_N$, let $\mu = N^{-1} \sum_i x_i$ and $m_k = N^{-1} \sum_i (x_i - \mu)^k$. We report the mean μ, variance m_2, skewness γ_1, and excess kurtosis γ_2, where

$$\gamma_1 = \frac{m_3}{(m_2 + \epsilon)^{3/2} + \epsilon}, \qquad \gamma_2 = \frac{m_4}{m_2^2 + \epsilon} - 3, \tag{A2}$$

with $\epsilon = 10^{-12}$. Histograms are normalized to unit mass.

Tone and color. The 32 tone descriptors comprise four moments of I, a 16-bin HSV-value histogram, seven luminance quantiles $(1, 5, 25, 50, 75, 95, 99\%)$, dark and bright clipping ratios at $1/255$ and $254/255$, 256-bin luminance entropy (Shannon 1948), and the mean and standard deviation of 16×16 local RMS contrast. The 60 color descriptors comprise moments of each RGB channel, 16-bin hue and saturation histograms, moments of CIELAB a^* and b^* (Luo 2023), the three pairwise RGB correlations, Hasler–Süsstrunk colorfulness (Hasler and Suesstrunk 2003), the mean sine and cosine of hue, and the mean and standard deviation of Lab chroma.

Texture and edges. The 38 texture descriptors include Gray-Level Co-occurrence Matrixs (GLCMs) (Haralick, Shanmugam, and Dinstein 1973), uniform Local Binary Patterns (LBPs) (Ojala, Pietikainen, and Maenpaa 2002), and local entropy. GLCMs use 32 quantized levels of I, distances $\{1, 2\}$, and orientations $\{0, \pi/4, \pi/2, 3\pi/4\}$; we retain the mean and standard deviation of contrast, energy, correlation, and homogeneity. Uniform LBPs contribute 10 bins at $(P, R) = (8, 1)$ on I and 18 bins at $(16, 2)$ on quantized Y, followed by the mean and standard deviation of radius-5 local entropy on Y. The 20 edge and sharpness descriptors comprise Canny densities (Canny 1986) at $\sigma \in \{1, 2, 4\}$, the gradient-magnitude mean, standard deviation, and maximum, four Histogram of Oriented Gradients (HOG) summaries (Dalal and Triggs 2005) (mean, standard deviation, maximum, and nonzero fraction), a magnitude-weighted 9-bin Sobel-orientation histogram, and Laplacian variance. The $\sigma = 1$ Canny, gradient, and HOG features use I, while the remaining descriptors use Y; HOG uses 9 orientations, 16×16 cells, 2×2 blocks, and L2-Hys normalization.

Frequency. The 21 frequency descriptors comprise an Fast Fourier Transform (FFT) high/low-energy ratio on I with cutoff $0.1 \min(H, W)$, normalized energy in four equal-width radial bands of the same spectrum, mean Gabor amplitudes (Gabor 1946) on I for four orientations and three frequencies $\{0.1, 0.2, 0.3\}$, and four power-spectrum summaries of Y. For the latter, let $P = |\mathcal{F}(Y)|^2$, let $r \in [0, 1]$ denote normalized radial frequency, and let $Z = \sum P$ over

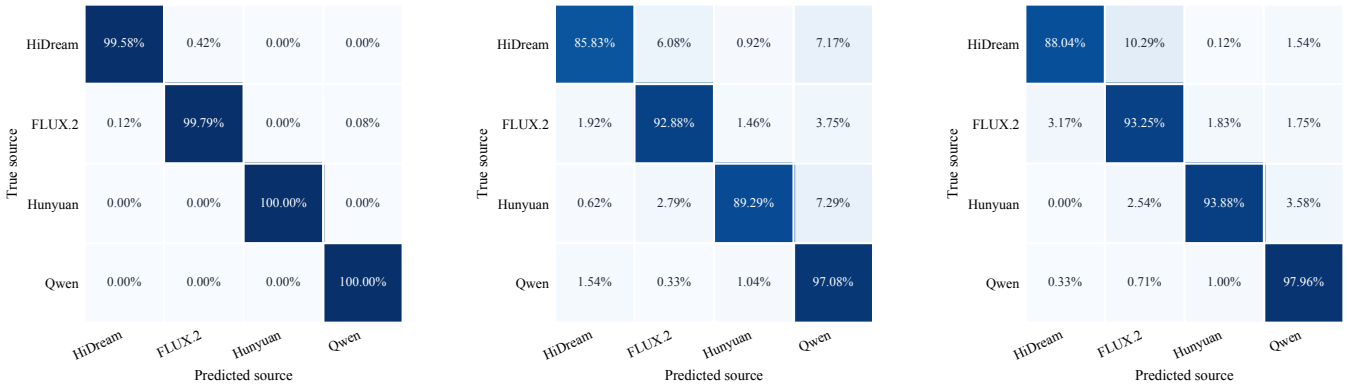

(a) All features. (b) Tone only. (c) Color only.

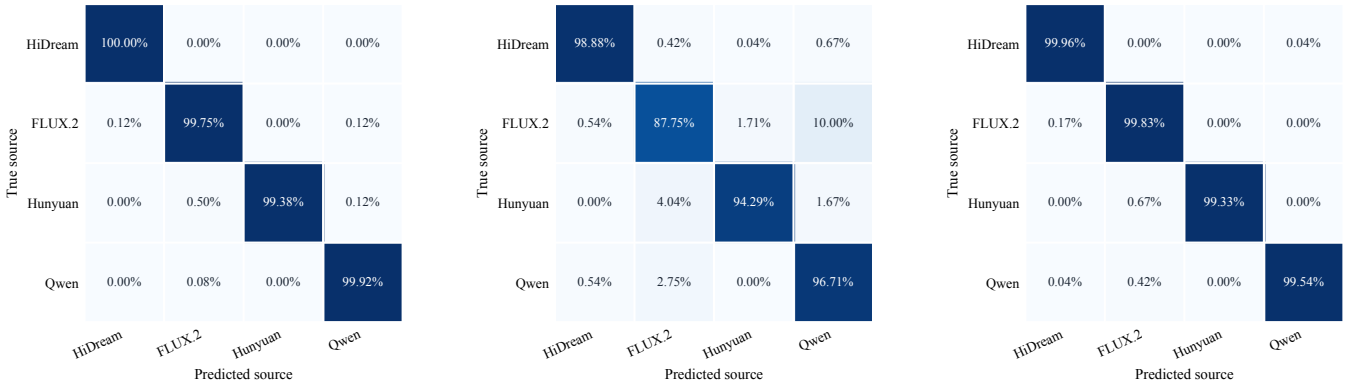

(d) Texture only. (e) Edge and sharpness only. (f) Frequency only.

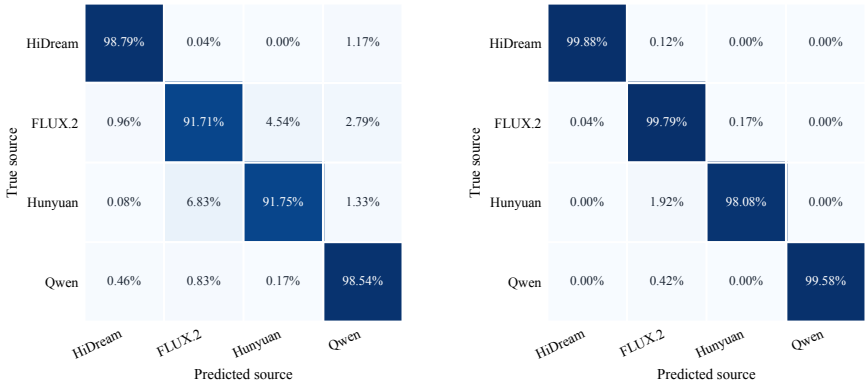

(g) Spatial organization only. (h) Residual and noise only.

Figure A3: **Source-classification confusion matrices for the full handcrafted feature bank and for each feature category alone.** Rows denote true source models and columns denote predicted sources. Each cell reports the row-normalized percentage, aggregated over three training seeds on the held-out ImageNet-derived test set.

non-DC coefficients:

$$c = \sum rP/Z,$$
$$s = \left[\sum (r-c)^2 P/Z\right]^{1/2}, \tag{A3}$$
$$f = \exp(\mathbb{E}[\log(P+\epsilon)])/(\mathbb{E}[P]+\epsilon).$$

These quantities are the spectral centroid, spread, and flatness. The spectral slope is fitted by linear regression of log mean power on log radial frequency over 32 radial bins, excluding the innermost bin.

Spatial organization and residual. The 40 spatial-organization descriptors combine morphology and coarse layout. An Otsu mask (Otsu 1979) of I provides mask occupancy, connected-component count, the mean and standard deviation of normalized component area, solidity, and circularity, and mean extent. A 3×3 grid contributes the luminance mean, luminance standard deviation, and Canny-$\sigma = 1$ density of every cell. We additionally compute mean absolute left–right and top–bottom reflection errors

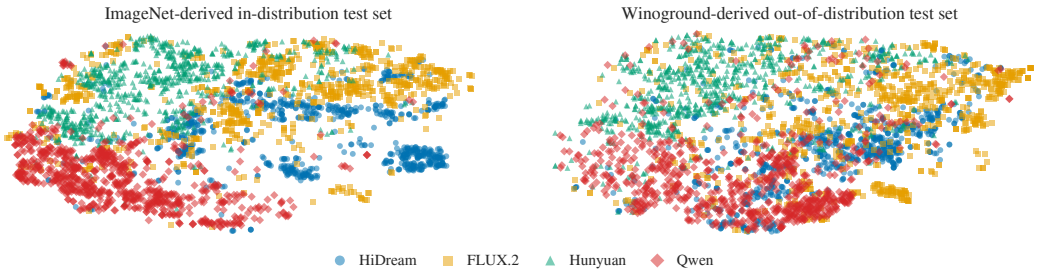

(a) Tone.

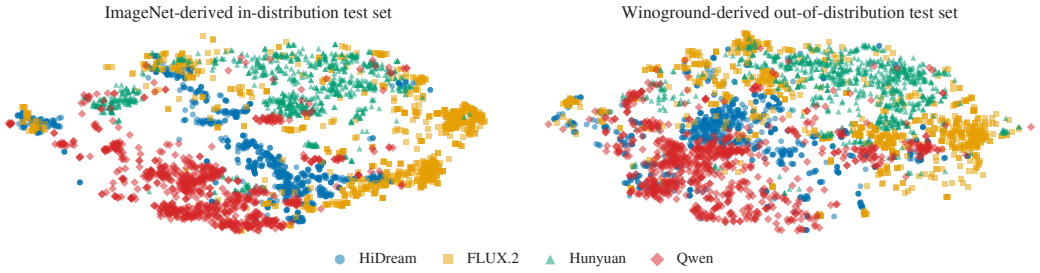

(b) Color.

Figure A4: **t-SNE visualization of each handcrafted feature category.** A joint embedding is fitted to balanced ID and OOD subsets and displayed separately per domain with shared coordinates. t-SNE is applied directly to the standardized features without PCA preprocessing. Continued in Fig. A4.

of Y and the luminance-weighted centroid in normalized image coordinates. The 14 residual and noise descriptors use $R = Y - \mathcal{G}_{\sigma=1} * Y$: four moments, the 50th, 90th, and 99th percentiles of $|R|$, one-pixel horizontal, vertical, and diagonal autocorrelations, and normalized energy in four radial FFT bands.

Implementation. We implement the extractor with NumPy (Harris et al. 2020), OpenCV (Bradski and Kaehler 2008), and scikit-image (Van der Walt et al. 2014), and the classifier in PyTorch (Paszke et al. 2019). Raw feature vectors are stored as `float32`. All 22,400 generated images are verified to be 1024×1024, and all extracted values are finite.

C.3 Handcrafted-Feature Classifier

Features follow the data split of Sec. A.2. Each feature is standardized with ImageNet-derived training statistics,

$$\widetilde{x}_{ij} = (x_{ij} - \mu_j^{\text{train}})/\sigma_j^{\text{train}}, \tag{A4}$$

where $\sigma_j^{\text{train}} < 10^{-6}$ is set to one; this affects only `residual_moment_mean`. The Multi-Layer Perceptron (MLP) is $d \to 256 \to 128 \to 4$, where d is the input dimension, with ReLU and dropout (Srivastava et al. 2014) of 0.1 after both hidden layers. We train with cross-entropy, Adam (Kingma and Ba 2015) at 10^{-4}, batch size 16, and 100 epochs, retaining the checkpoint with the highest validation accuracy. Hyperparameters are fixed across configurations. Reported accuracies are means over seeds 42, 43, and 44.

C.4 Feature-Category Analysis

We evaluate each feature category from two complementary perspectives. For category sufficiency, we train the same MLP using only the descriptors in that category. For grouped permutation importance (Breiman 2001; Gregorutti, Michel, and Saint-Pierre 2015), we evaluate the trained full-bank classifier after applying a common random row permutation to all descriptors in one category while leaving the other categories unchanged; this preserves dependencies within the permuted category but breaks its association with the image and its source label. Permutations are performed separately for each dataset and training seed. Both analyses use the same train-only normalization, data splits, architecture, and checkpoint-selection procedure described above.

C.5 Confusion Matrices and t-SNE Embeddings

Figure A3 reports the row-normalized confusion matrices for the complete feature bank and each feature category alone. The complete bank achieves 99.84% ID test accuracy, while texture, frequency, and residual and noise features alone reach 99.76%, 99.67%, and 99.33%, respectively. Tone, color, edge and sharpness, and spatial organization reach 91.27–95.20%. Class-wise performance varies with the feature category. HiDream has the highest recall for edge and sharpness (98.88%) and spatial organization (98.79%), computed by pooling confusion counts across the three training seeds. Multiple feature categories therefore independently retain strong source information.

We visualize each feature category with t-SNE (Van der Maaten and Hinton 2008). For every category, we standardize the corresponding feature dimensions and apply t-SNE

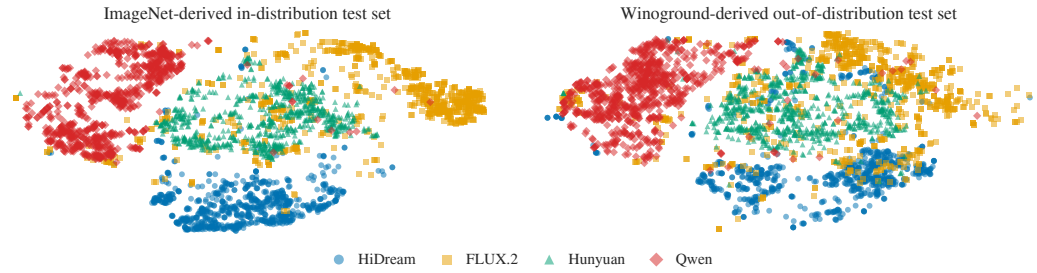

(c) Texture.

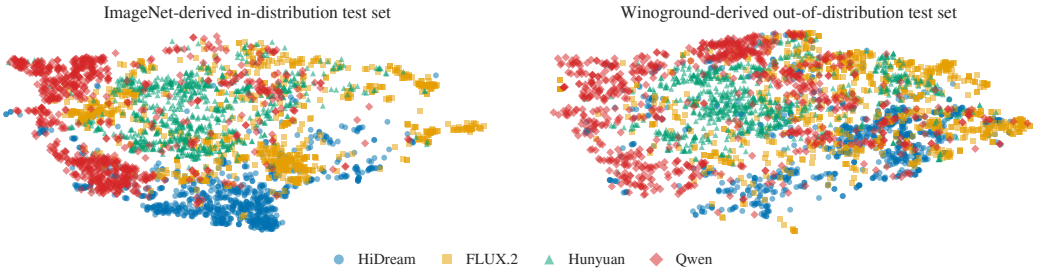

(d) Edge and sharpness.

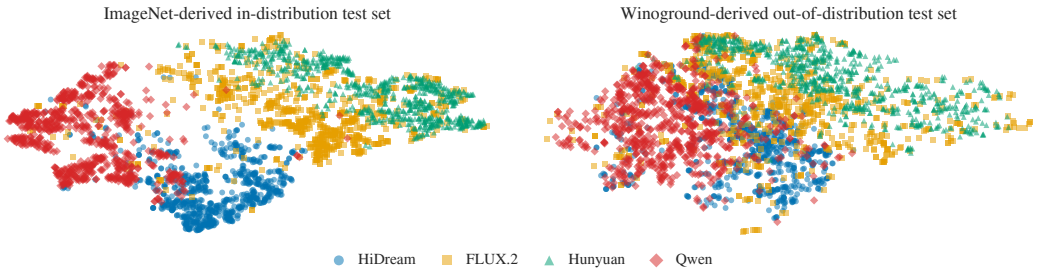

(e) Frequency.

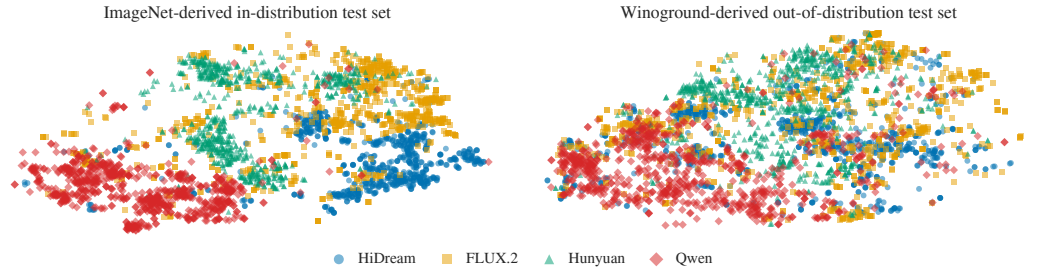

(f) Spatial organization.

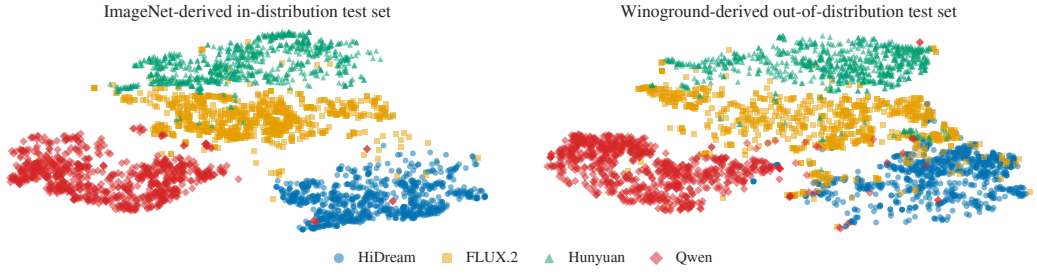

(g) Residual and noise.

Figure A4: **t-SNE visualization of each handcrafted feature category (continued).** Texture is the only category whose embedding separates all four source models in both domains.

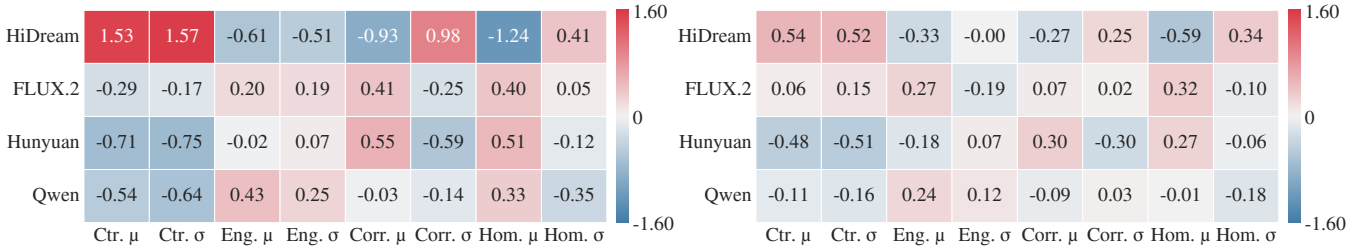

(a) GLCM statistics on the ImageNet-derived set.

(b) GLCM statistics on the Winoground-derived set.

Figure A5: **Prompt-centered GLCM statistics.** Cells report source-wise mean prompt-centered z-scores for the mean and standard deviation of GLCM contrast, energy, correlation, and homogeneity. Source-specific signed profiles remain visible on both prompt distributions.

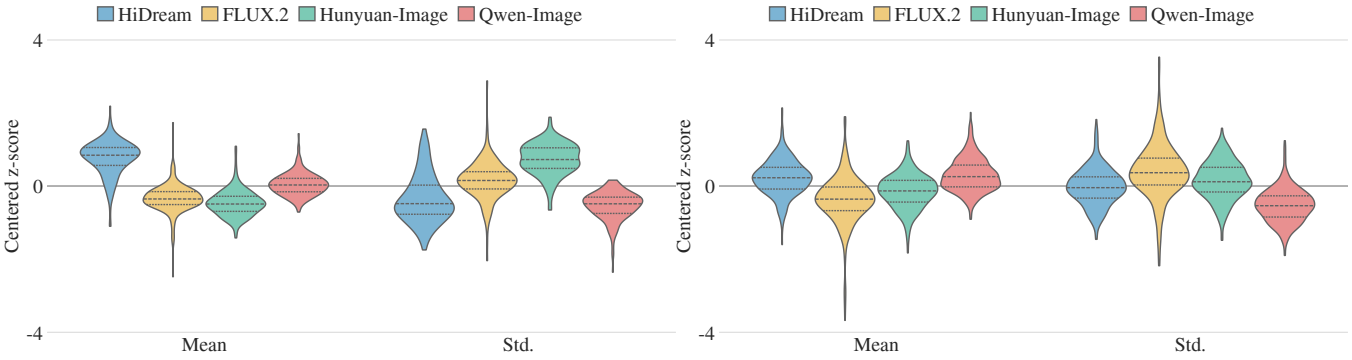

(a) Local entropy on the ImageNet-derived set.

(b) Local entropy on the Winoground-derived set.

Figure A6: **Prompt-centered local-entropy distributions.** Violin plots show the distributions of the per-image local-entropy mean and standard deviation after matched-prompt centering. HiDream's entropy mean is above the four-source average in both domains. Its entropy standard deviation is below this average on ImageNet and close to it on Winoground.

directly, without a preceding PCA projection. A single joint embedding is fitted to balanced ImageNet-derived ID and Winoground-derived OOD subsets, and the two domains are displayed separately with the same embedding coordinates, which makes changes in source separation across prompt distributions directly visible. As shown in Figs. A4 and A4, texture features form a clear four-source organization in both domains. Frequency and residual-and-noise features retain partial source structure on the ImageNet-derived set, but several sources overlap substantially, and the overlap grows on Winoground. Tone, color, edge and sharpness, and spatial-organization embeddings overlap extensively and do not form four distinct clusters.

C.6 Texture Statistics

We examine the constituent statistics of the texture category in greater detail. Before aggregation, each feature is centered across the four source models within every matched seed-prompt group, and the four seeds of each prompt are then averaged for each source. This removes the component shared by images generated from the same prompt and emphasizes systematic source-conditioned offsets rather than prompt-content differences.

As shown in Fig. A5, the GLCM statistics reveal distinct co-occurrence profiles for the four source models. HiDream and HunyuanImage tend to occupy opposite directions on several contrast- and energy-related statistics, while FLUX.2 and Qwen-Image exhibit different correlation and homogeneity offsets. The magnitudes change between ImageNet and Winoground, but the signed source-dependent profiles remain visible in both domains, so the signal captured by GLCM features is not reducible to a single notion of contrast or smoothness.

Fig. A6 shows complementary differences in local entropy. HiDream's local-entropy mean is above the four-source average on both ImageNet and Winoground, with average prompt-centered z-scores of +0.765 and +0.227, respectively. Its local-entropy standard deviation is below this average on ImageNet (-0.343) and close to it on Winoground (-0.022). Qwen-Image also has positive entropy-mean offsets and negative entropy-standard-deviation offsets on both prompt distributions.

The uniform LBP histograms in Fig. A7 provide a more localized view of texture. At both $(P, R) = (8, 1)$ and $(P, R) = (16, 2)$, the four source models assign systematically different probability mass to specific pattern bins, with the largest deviations concentrated in a small subset of bins rather than spread over the histogram. Several prominent offsets, including the distinctive HiDream and FLUX.2 profiles, persist under the Winoground prompt shift even though

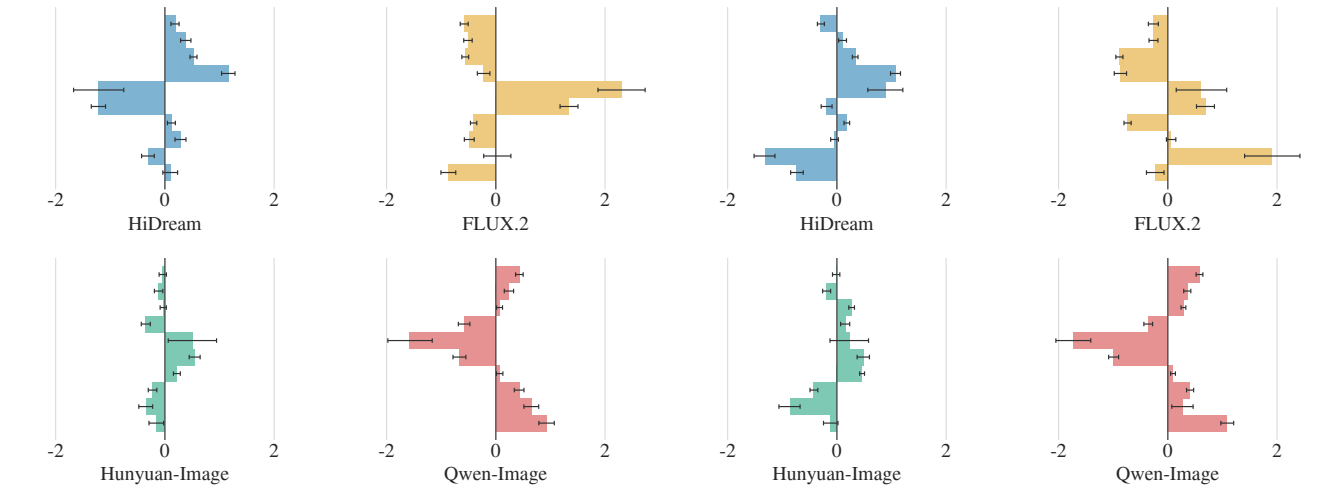

(a) ImageNet-derived set, $(P, R) = (8, 1)$.

(b) Winoground-derived set, $(P, R) = (8, 1)$.

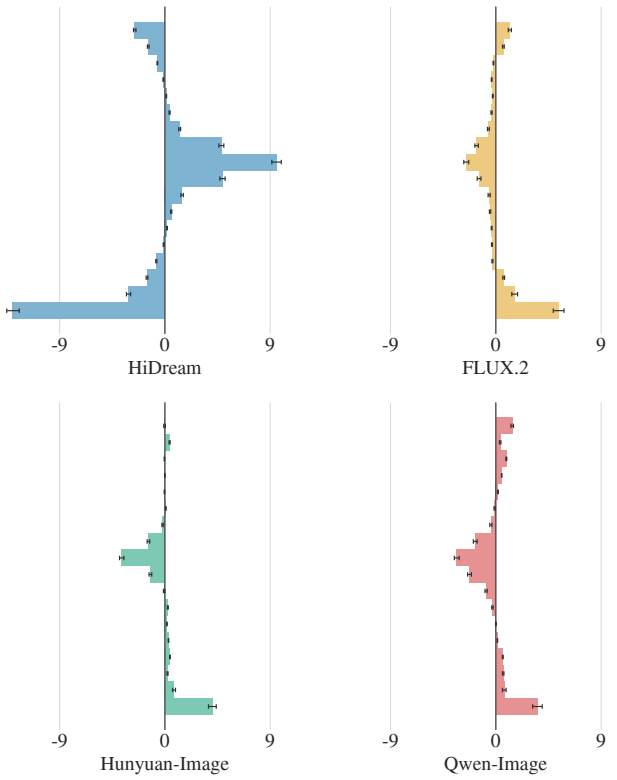

(c) ImageNet-derived set, $(P, R) = (16, 2)$.

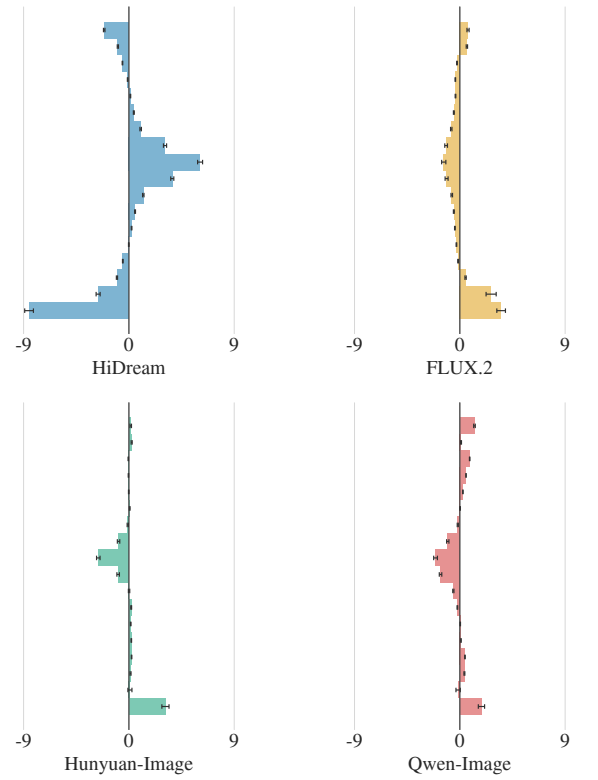

(d) Winoground-derived set, $(P, R) = (16, 2)$.

Figure A7: **Prompt-centered uniform-LBP profiles.** Within each model panel, horizontal bars are ordered by uniform-LBP bin index from top to bottom: bins 0–9 for $(P, R) = (8, 1)$ and bins 0–17 for $(P, R) = (16, 2)$, with the final bin collecting non-uniform patterns. For every seed-prompt group and bin, we subtract the mean bin probability across the four source models and average the centered values over the four seeds of each prompt. Bar lengths report the mean prompt-centered bin-probability difference across prompts in percentage points, with zero denoting the matched four-source mean.

their amplitudes change. Together, the GLCM, local-entropy, and LBP results show that texture-based source classification combines several forms of local structure rather than relying on one scalar texture statistic.

D Mid-Level Visual Representations

D.1 Representation Extraction

We extract foreground masks, metric depth, and albedo from every generated image. For foreground masks, we use text-

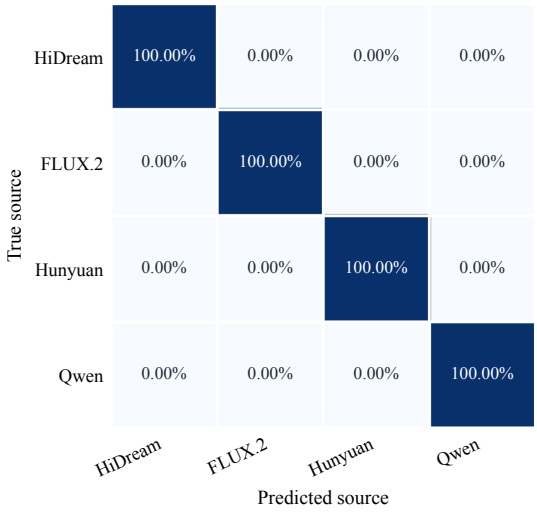

(a) Original image.

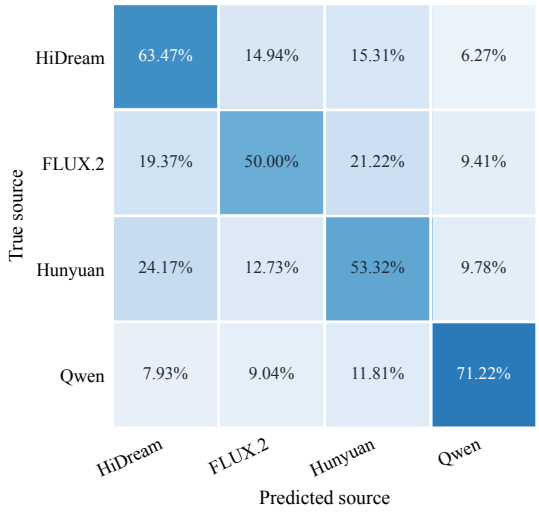

(b) Segmentation mask.

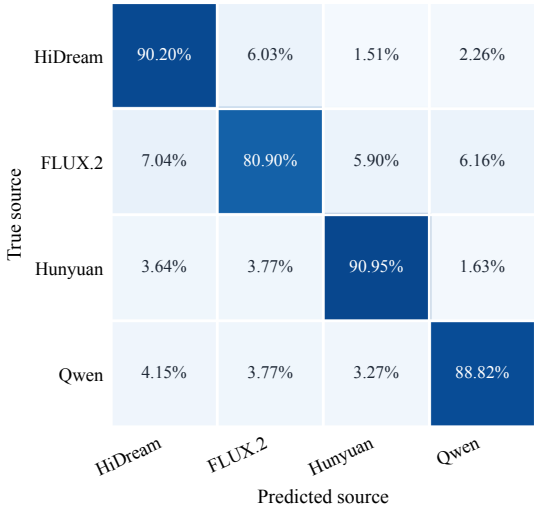

(c) Depth map.

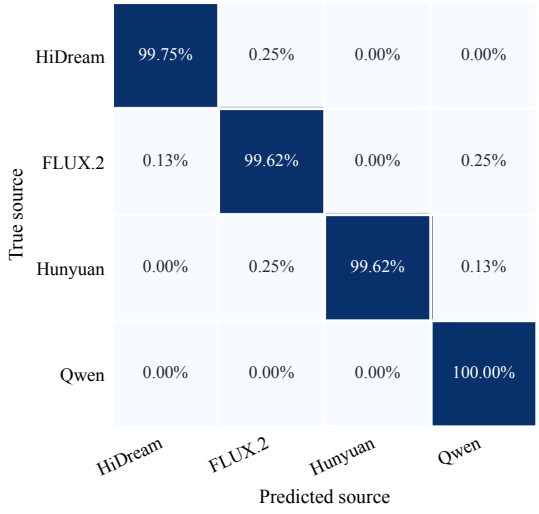

(d) Albedo map.

Figure A8: **Source-classification confusion matrices for the mid-level visual representations.** All four classifiers are ResNet-50. Rows denote true source models and columns denote predicted sources. Each cell reports the row-normalized percentage.

guided segmentation with SAM 3.1 (Carion et al. 2026). On ImageNet-derived images, we query SAM 3.1 once with the corresponding ImageNet class label and take the pixel-wise union of all returned instance masks. On Winoground-derived images, we query SAM 3.1 separately with the two extracted object names and take the pixelwise union of the masks returned by both queries. In both cases, the foreground mask is defined by the alpha channel of the resulting cutout.

We apply the off-the-shelf MoGe-2 (Wang et al. 2025a) model to each image and retain its per-pixel metric-depth prediction. The depth representation given to the source classifier, however, is the rendered relative-depth image rather than the raw metric-depth array. For each image i, we convert every valid metric depth $z_i(p) > 0$ to disparity $u_i(p) = 1/z_i(p)$ and compute the 0.001 and 0.99 disparity quantiles, $q^{(i)}_{0.001}$ and $q^{(i)}_{0.99}$, over all valid pixels. The relative depth is

$$r_i(p) = \mathrm{clip}\left(1 - \frac{u_i(p) - q^{(i)}_{0.001}}{q^{(i)}_{0.99} - q^{(i)}_{0.001}}, 0, 1\right). \tag{A5}$$

We map r_i to RGB with the MoGe-2 depth colormap and use the rendered image as the classifier input. The classification representation is thus normalized independently for every generated image and does not encode absolute metric scale; the raw metric-depth values are used only for the metric-depth analysis in Sec. D.2.

We additionally apply the off-the-shelf DiffusionRenderer (Liang et al. 2025a) inverse-rendering model and retain its estimated base-color pass as the albedo map.

D.2 Layout, Depth, and Albedo Statistics

Average object layouts. We first remove masks whose foreground occupies less than 5% or more than 95% of the image. To preserve the matched comparison, if any source model fails this criterion for a given prompt and seed, we remove the corresponding images from all four sources, leaving 2,612 matched images per source model. For each retained 1024×1024 mask M_i, we approximate its signed distance field in pixel units as

$$d_i(p) = \mathrm{EDT}(M_i)(p) - \mathrm{EDT}(1 - M_i)(p), \tag{A6}$$

where EDT is the Euclidean distance transform, so that $d_i(p)$ is positive inside the foreground and negative outside it. For source model s, the average layout plotted in Fig. 6 is

$$\bar{d}_s(p) = \frac{1}{N_s} \sum_{i=1}^{N_s} d_i(p), \tag{A7}$$

computed directly in the original image coordinates without spatial alignment. We display $\bar{d}_s$ together with its zero-level contour; the reported mean-layout area is the fraction of pixels satisfying $\bar{d}_s(p) \geq 0$.

Metric depth of generated objects. For image i, we define the object region as the intersection of the SAM foreground and the pixels with finite MoGe-2 depth,

$$\Omega_i = \{p : M_i(p) = 1\} \cap \{p : z_i(p) \text{ is finite}\}, \tag{A8}$$

where $z_i(p)$ is the predicted metric depth in meters. The near- and far-depth summaries plotted in Fig. 7 are the 5th and 95th percentiles of $\{z_i(p) : p \in \Omega_i\}$, shown on a logarithmic depth axis.

Fine-scale variation ratio. We compute this quantity from the DiffusionRenderer albedo estimate inside the retained SAM foreground. Seven matched groups are excluded because at least one source has insufficient valid pixel pairs at a sampled distance. We remove all four source images in each affected group, leaving 2,605 images per source for this analysis. To reduce boundary artifacts, we erode each foreground by 5 Euclidean pixels and denote the remaining region by Ω_i. Let

$$L_i = \sqrt{(h_i - 1)^2 + (w_i - 1)^2} \tag{A9}$$

be the diagonal of the axis-aligned bounding box of Ω_i. For each normalized distance

$$r \in \{0.005, 0.01, 0.02, 0.04, 0.08, 0.16, 0.24, 0.32, 0.40\}, \tag{A10}$$

we sample 1,024 pixel pairs (p, q) within Ω_i whose Euclidean separation is approximately rL_i. After converting the base-color values to CIELAB, we define the raw albedo variogram as

$$D_i(r) = \underset{\substack{p,q \in \Omega_i \\ \|p-q\|_2 / L_i \approx r}}{\mathrm{median}} \Delta E_{00}(p, q), \tag{A11}$$

where ΔE_{00} is the CIEDE2000 color difference. We fit a nondecreasing curve $\widetilde{D}_i(r)$ to these values with weighted isotonic regression, using the number of valid pairs as the weight. The fine-scale variation ratio plotted in Fig. 8 is

$$R_{\mathrm{fine},i} = \frac{\widetilde{D}_i(0.005)}{\widetilde{D}_i(0.40)}, \tag{A12}$$

which measures how much of an object's long-range albedo difference is already present at a separation equal to 0.5% of the object scale.

Representation	Classifier	Acc. (ID)	Acc. (OOD)
Segmentation mask	ResNet-50	59.50	46.73
	Two-layer CNN	56.78	42.29
	DINOv2 ViT-B/14	50.46	36.17
	CLIP ViT-B/32	44.56	32.46
Depth map	ResNet-50	87.72	70.50
	Two-layer CNN	76.66	56.67
	DINOv2 ViT-B/14	69.22	56.22
	CLIP ViT-B/32	60.96	45.95
Albedo map	ResNet-50	99.75	97.73
	Two-layer CNN	95.41	88.58
	DINOv2 ViT-B/14	86.81	80.41
	CLIP ViT-B/32	81.66	72.08

Table A7: **Source classification from mid-level visual representations across classifier architectures.** ID denotes evaluation on held-out ImageNet-derived prompts and OOD denotes evaluation on Winoground-derived prompts without target-domain training. All values are accuracies in percent; chance is 25%.

D.3 Confusion Matrices

Figure A8 compares the confusion matrices obtained by ResNet-50 from RGB images and the three estimated mid-level representations. RGB classification reaches 100.00% ID evaluation accuracy. Albedo achieves 99.75%, with eight misclassified images among 3,200, and depth achieves 87.72%. Segmentation masks reach 59.50% on the 2,168 retained ID images, with Qwen-Image having the highest class-wise recall. These results demonstrate source information in object layout, geometry, and estimated surface appearance.

D.4 Classification Across Architectures

Table A7 reports results for all combinations of the three representations and the four classifier architectures. The ordering of the representations is consistent across every classifier and both prompt distributions: albedo is most predictive, depth intermediate, and segmentation masks least predictive. ResNet-50 gives the strongest result for each representation, but the two-layer CNN and the linear classifiers on frozen DINOv2 and CLIP features also remain above the 25% chance level. Accuracy decreases on Winoground-derived prompts, while the ordering across representations is preserved. Albedo remains highly discriminative, reaching 97.73% OOD accuracy with ResNet-50 and 72.08% with a linear classifier on frozen CLIP features.

E High-Level Semantic Realization

E.1 Captioning Setup

We use Qwen3-VL-8B-Instruct (Bai et al. 2025), InternVL3.5-8B (Wang et al. 2025b), and GLM-4.1V-9B-Thinking (Hong et al. 2025) as captioners. Each captioner processes all 16,000 ImageNet-derived and 6,400

Winoground-derived images independently under each of the two prompts below, yielding 44,800 captions per captioner. The VLM input contains only the RGB image and the instruction; source labels, filenames, paths, and the original generation prompt are not provided. We use bfloat16 inference, greedy decoding without sampling, and a maximum of 192 new tokens. For the GLM thinking checkpoint, we suppress chain-of-thought generation and retain only the final answer text. Example captions are shown in Tab. A9.

Style-description prompt

```
Describe the image's overall visual style
and atmosphere in one or two sentences.
Focus on how the image looks and feels as a
whole rather than listing the objects or
actions it contains. Use only directly
observable qualities.
```

Background-description prompt

```
Describe only the visible background of
this image in one or two factual sentences.
Focus on the setting and on what appears
behind or around the main subject,
including the background's overall visual
appearance. Do not describe or identify the
foreground subject. If the background is
absent, blurred, or unclear, state that
briefly instead of inferring details.
```

E.2 Caption Classification

The source-classification task has the same four classes and the same image-level splits as Sec. A.2; captions never determine the split.

Text preprocessing. Style and background captions are analyzed separately for every captioner; we do not concatenate captions across prompts or captioners. We remove leading and trailing whitespace and, for GLM, its answer-formatting markers. We otherwise preserve the generated text: no manual stop-word removal, stemming, lemmatization, or punctuation filtering is applied. Lowercasing is performed by the Term Frequency-Inverse Document Frequency (TF-IDF) vectorizers, with Unicode accent normalization for the word stream.

Word- and character-level features. We concatenate two sparse TF-IDF representations (Salton and Buckley 1988) implemented with scikit-learn (Pedregosa et al. 2011). The word branch uses unigrams and bigrams, while the character branch uses within-word character 3- to 5-grams. Both branches use sublinear term frequency, smoothed inverse document frequency, L_2 normalization, a minimum document frequency of two, and at most 100,000 features; the word branch additionally applies Unicode accent stripping. For each captioner-prompt condition, both vocabularies and all inverse-document-frequency statistics are fitted only on the ImageNet-derived training captions and then held fixed for validation, held-out ID testing, and Winoground OOD testing.

Logistic regression. For each of the six captioner-prompt conditions, we fit a separate four-class, L_2-regularized multinomial logistic-regression classifier with the L-BFGS solver (Liu and Nocedal 1989), an iteration limit of 1,000, and random seed 0, without class reweighting. The training set contains 12,800 captions; the held-out ID, and OOD sets contain 3,200, and 6,400 captions, respectively. Accuracy is the primary metric, with chance at 25%. We select the inverse regularization strength $C \in \{0.1, 1, 10\}$ by ImageNet-derived validation accuracy, independently for each captioner and prompt, and evaluate the selected classifier unchanged on the held-out ID and OOD sets. The selected value is $C = 1$ for both Qwen3-VL conditions, both InternVL3.5 conditions, and GLM-4.1V background, and $C = 10$ for GLM-4.1V style.

E.3 Confusion Matrices

Figure A9 reports the row-normalized confusion matrices of the combined word- and character-level TF-IDF classifiers on the held-out ImageNet-derived test set. All six matrices are diagonally dominant. On the ID test set, style and background descriptions yield accuracies of 65.62% and 71.12% for Qwen3-VL, 66.06% and 70.94% for InternVL3.5, and 68.78% and 71.78% for GLM-4.1V, respectively. Background descriptions yield higher accuracy for every captioner.

Across the six combinations of captioner and description prompt, Qwen-Image has the highest recall (77.88–86.75%) and FLUX.2 has the lowest (52.25–64.00%). The largest off-diagonal entry for FLUX.2 is assignment to HunyuanImage in four conditions and to Qwen-Image in two. The dominant confusion direction therefore varies with the captioner and description prompt.

E.4 Model-Specific Preferences on Style

To interpret the lexical evidence underlying style-based source classification, we fit an auxiliary word-level TF-IDF logistic-regression classifier for each captioner and report in Tab. A8 the five terms with the largest positive coefficient for each source model; the corresponding results for background descriptions are in Tab. 7. The table retains both unigrams and bigrams, consistent with the word-level representation used in the classification experiments. Coefficients indicate predictive associations and should not be interpreted causally.

Several associations recur across all three captioners. FLUX.2 is associated with *smooth*, HunyuanImage with *dreamy*, and Qwen-Image with *dark* and *green*. HiDream's top-five lists vary across captioners, with terms including *soft natural*, *window*, *brick*, *lush*, *homey*, and *cozy*. These coefficients describe source-associated language patterns, with agreement depending on the model and captioner.

F Qualitative Examples

Figure A10 shows images generated from ImageNet-derived and Winoground-derived prompts, Fig. A11 shows the corruptions of Sec. B.3, Fig. A12 visualizes the mid-level representations of Sec. D.1, and Tab. A9 lists captions produced under the two prompts of Sec. E.1.

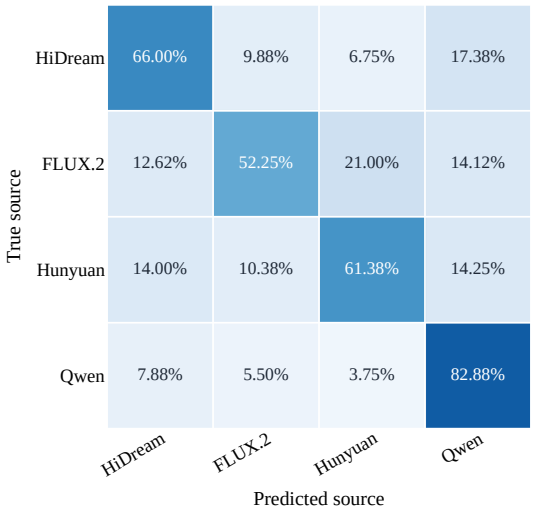

(a) Qwen3-VL-8B, style.

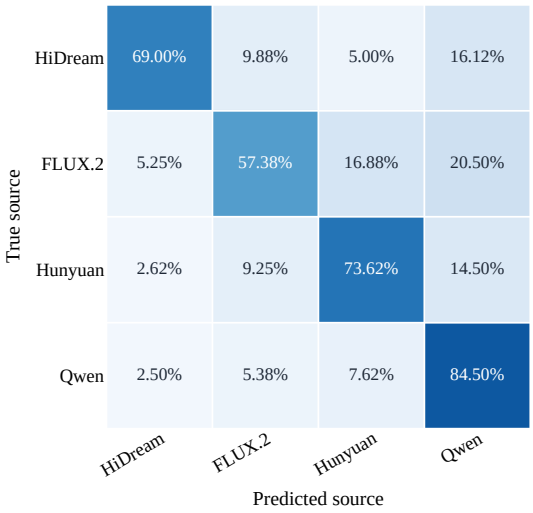

(b) Qwen3-VL-8B, background.

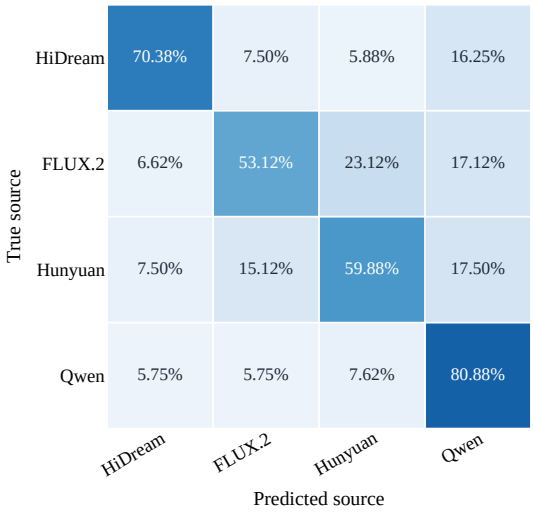

(c) InternVL3.5-8B, style.

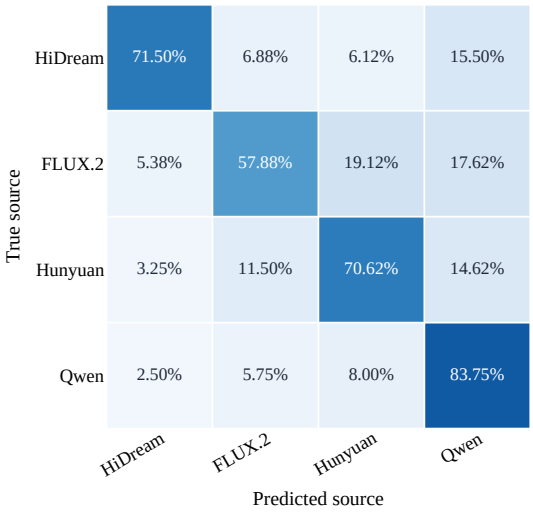

(d) InternVL3.5-8B, background.

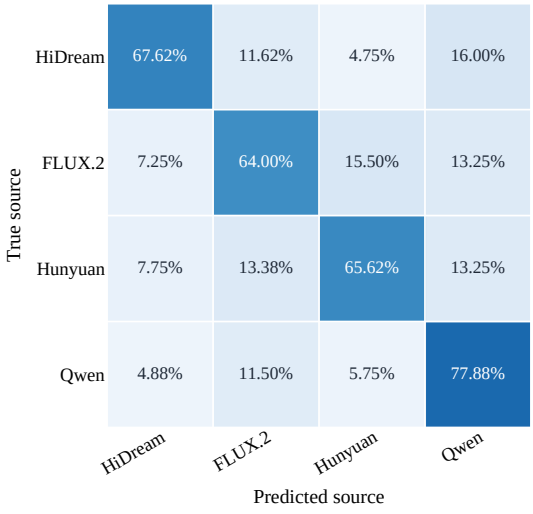

(e) GLM-4.1V-9B, style.

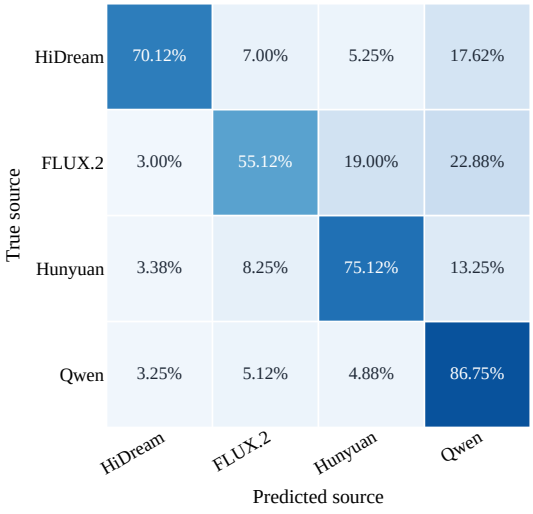

(f) GLM-4.1V-9B, background.

Figure A9: **Caption-based source-classification confusion matrices for style and background descriptions from three VLM captioners.** Each classifier uses combined word- and character-level TF-IDF features and is evaluated on the held-out ImageNet-derived test set. Rows denote true source models and columns denote predicted sources. Each cell reports the row-normalized percentage.

HiDream	FLUX.2	Hunyuan-Image	Qwen-Image
soft natural	**smooth**	hazy	**dark**
window	minimalist	dusty	**green**
brick	overcast	**dreamy**	moody
natural light	golden	bokeh	damp
naturalistic	clean	dust	verdant
background of	image features	warm	**dark**
wooden	minimalist	**dreamy**	**green**
window	simple	sunlight	**green** background
lush	features	warm soft	against **dark**
greenery	**smooth**	gentle	muted
natural and	has warm	**dreamy**	moody
lush	**smooth**	dusty	has moody
homey	warm soft	sunlit	**green**
cozy	clean	warm	intimate
brick	overcast	hazy	**dark**

Table A8: **Model-specific style preferences revealed by VLM-generated captions.** We list the top-5 predictive word n-grams from word-level TF-IDF logistic-regression classifiers for style descriptions. The three row blocks correspond, from top to bottom, to Qwen3-VL-8B, InternVL3.5-8B, and GLM-4.1V-9B. Bold words are unigrams that rank among the top terms for the same source model under all three captioners.

G Limitations and Future Work

Our study is intentionally scoped to four open-source frontier Text-to-Image (T2I) models, two English prompt families, and one default generation configuration per model. The resulting accuracies therefore establish strong closed-set distinguishability within this roster, but do not guarantee attribution to proprietary models or future checkpoints. Moreover, each source label jointly reflects model weights, architecture, training data, sampler, decoder, and any model-specific post-processing; our analyses reveal product-level visual idiosyncrasies but do not causally localize them to a particular component.

Future work should disentangle these factors through controlled comparisons across checkpoints, model families, samplers, guidance scales, inference steps, resolutions, and shared decoding pipelines, while extending evaluation to open-set and continually updated model rosters. In addition, the mid-level maps and high-level descriptions are estimated by pretrained segmentation, depth, inverse-rendering, and captioning models and may inherit their biases or errors. Replication with alternative extractors, human annotations, and direct interventions on the identified visual properties would help distinguish stable generator behavior from measurement artifacts and clarify which idiosyncrasies causally support human and machine source classification.



HiDream FLUX.2 Hunyuan Qwen

(a) ImageNet-derived prompts. Above: “a photo of tench, which is also known as Tinca tinca.” Below: “a photo of black swan, which is also known as Cygnus atratus.”



HiDream FLUX.2 Hunyuan Qwen

(b) Winoground-derived prompts. Above: “an old person kisses a young person.” Below: “the happy one is on the left and the upset one is on the right.”

Figure A10: **Examples of generated images.** Each row shows the four source models’ outputs for the same prompt.

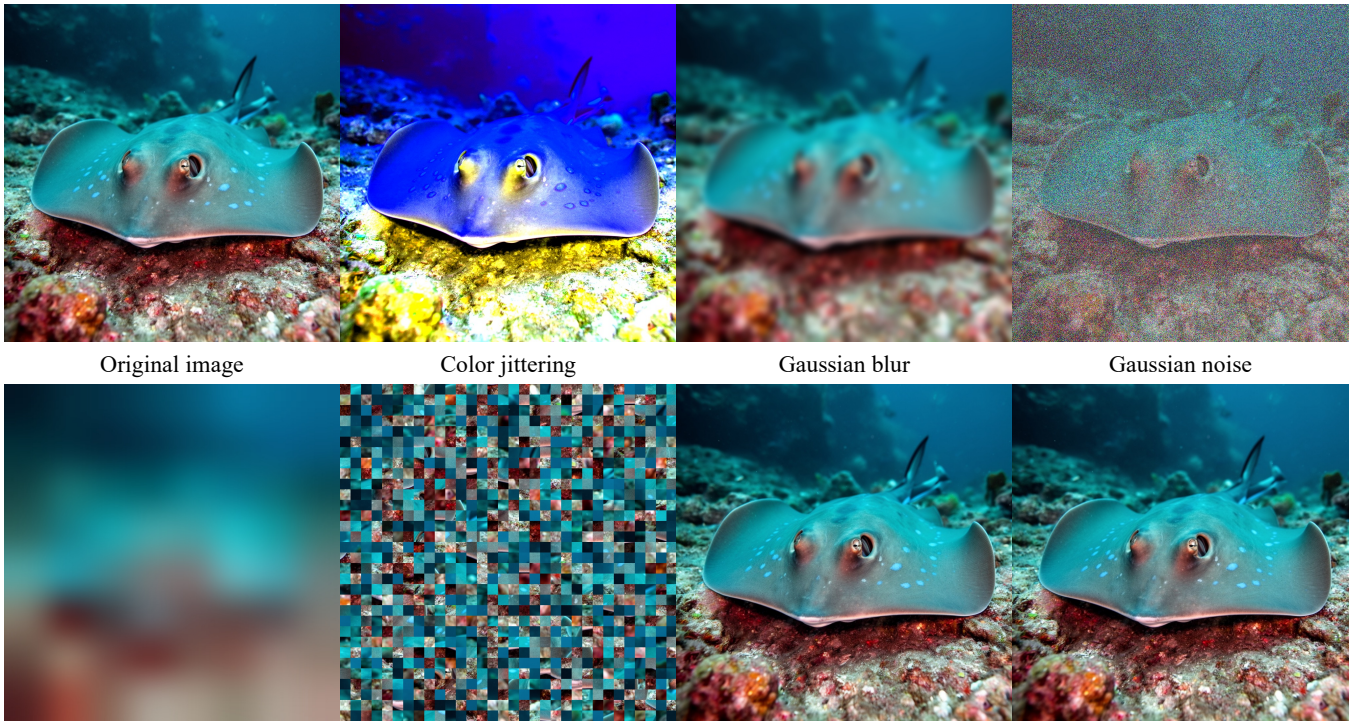

Figure A11: **Examples of corrupted images.** Each column applies one corruption setting of Tab. A3 to the same image.

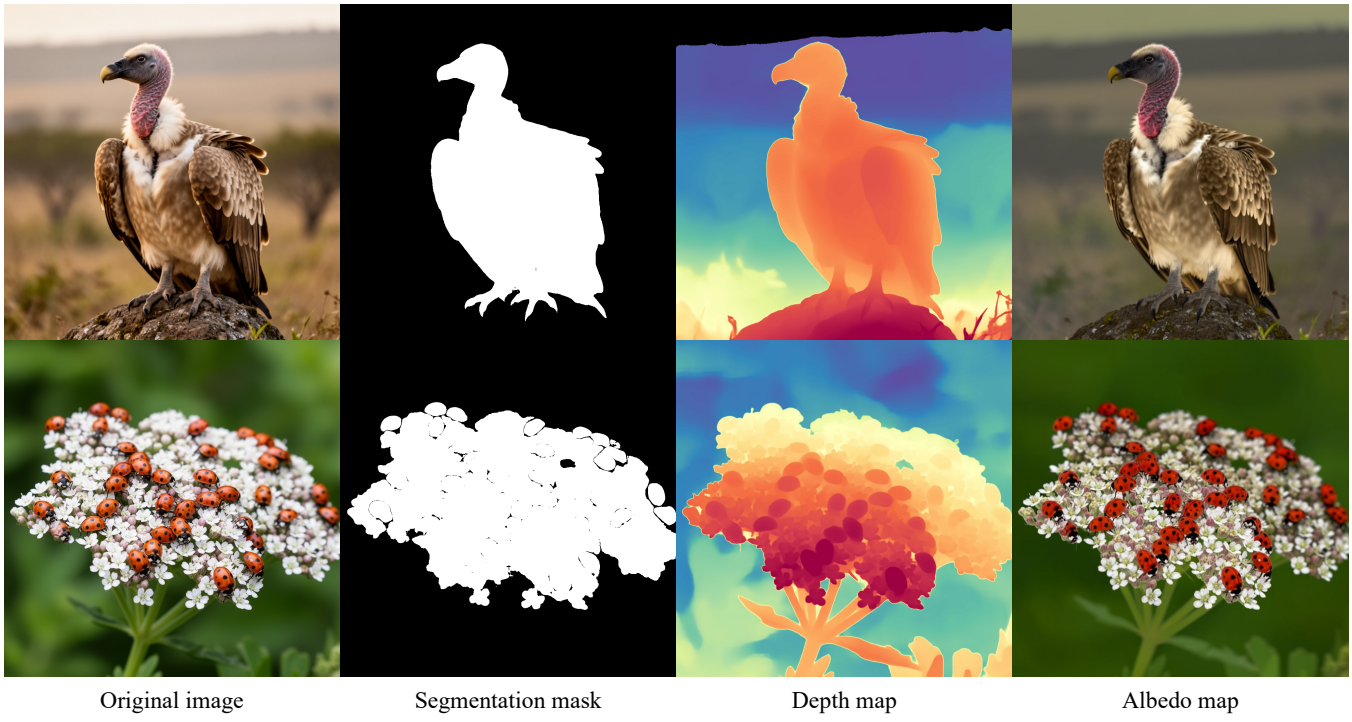

Figure A12: **Examples of mid-level representations.** Segmentation masks, relative-depth renderings, and albedo maps for images generated from ImageNet-derived prompts (above) and Winoground-derived prompts (below).

Image	Caption
	**Background-description prompt (Qwen3-VL)** The background consists of out-of-focus green foliage with soft, blurred patches of pinkish-purple flowers visible in the upper right and left corners, creating a natural, garden-like setting with shallow depth of field. **Style-description prompt (Qwen3-VL)** The image presents a vivid, sharply focused macro view of a strikingly patterned spider resting on a textured green leaf, with a soft, blurred natural background that enhances the subject's bold black-and-yellow contrast and creates a sense of quiet, detailed observation in a lush, living environment.
	**Background-description prompt (InternVL3.5)** The background of the image is a sandy seabed with some scattered debris and small particles, giving it a natural underwater appearance. The area behind the main subject is slightly blurred, emphasizing the foreground. **Style-description prompt (InternVL3.5)** The image has a vibrant and intense visual style, with the bright colors of the lobster contrasting against the dark, murky background, creating a dramatic and somewhat menacing atmosphere. The close-up perspective adds a sense of immediacy and detail, drawing the viewer into the scene.
	**Background-description prompt (GLM-4.1V)** The background shows a large crowd of people gathered in what appears to be an outdoor setting with dim lighting and string lights visible. There's a hazy atmosphere with some light particles or dust in the air, suggesting a festival or event environment. **Style-description prompt (GLM-4.1V)** The image has a cinematic, atmospheric style with a soft focus and warm tones, creating a moody, almost dreamlike quality. The scene feels dynamic and lively, with a sense of energy from the blurred crowd and particles in the air, while the dog stands as a clear focal point with a calm, dignified presence.

Table A9: **Examples of captions under the background and style prompts.** Captions are reproduced from the saved experimental outputs, with only typographic punctuation normalized.