

000 RAREFLOW: PHYSICS-AWARE FLOW-MATCHING 001 FOR CROSS-SENSOR SUPER-RESOLUTION OF RARE- 002 EARTH FEATURES 003

006 **Anonymous authors**

007 Paper under double-blind review

011 ABSTRACT

013 Super-resolution (SR) for remote sensing imagery often fails under out-of-
014 distribution (OOD) conditions, such as rare geomorphic features captured by di-
015 verse sensors, producing visually plausible but physically inaccurate results. We
016 present RareFlow, a physics-aware SR framework designed for OOD robustness.
017 RareFlow’s core is a dual-conditioning architecture. A Gated ControlNet pre-
018 preserves fine-grained geometric fidelity from the low-resolution input, while textual
019 prompts provide semantic guidance for synthesizing complex features. To ensure
020 physically sound outputs, we introduce a multifaceted loss function that enforces
021 both spectral and radiometric consistency with sensor properties. Furthermore,
022 the framework quantifies its own predictive uncertainty by employing a stochastic
023 forward pass approach; the resulting output variance directly identifies unfamiliar
024 inputs, mitigating feature hallucination. We validate RareFlow on a new, curated
025 benchmark of multi-sensor satellite imagery. In blind evaluations, geophysical
026 experts rated our model’s outputs as approaching the fidelity of ground truth im-
027 agery, significantly outperforming state-of-the-art baselines. This qualitative su-
028 periority is corroborated by quantitative gains in perceptual metrics, including a
029 nearly 40% reduction in FID. RareFlow provides a robust framework for high-
030 fidelity synthesis in data-scarce scientific domains and offers a new paradigm for
031 controlled generation under severe domain shift.

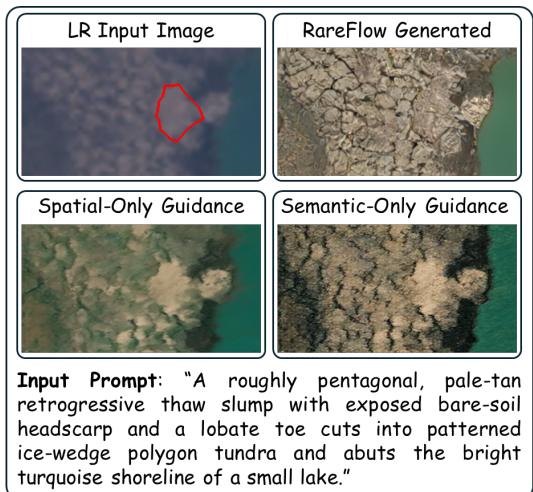
033 1 INTRODUCTION

035 Monitoring rapid, small-scale environmental change requires imagery that is both high spatial res-
036 olution (for morphology) and high temporal frequency (for dynamics) Qi et al. (2025); Vu et al.
037 (2025). Public constellations such as Sentinel-2 deliver near-global coverage with 10 m pixels and
038 5-day revisit, but lack the fine spatial detail needed to resolve many geomorphic features; very-
039 high-resolution (VHR) commercial sensors provide sub-meter detail at lower cadence and are often
040 expensive. European Space Agency (ESA) (a;b) This mismatch creates a persistent spatiotemporal
041 gap for near-real-time environmental monitoring. Cross-sensor super-resolution (SR) has significant
042 scientific value, as it promises to synthesize VHR content from publicly available lower-resolution
043 satellite imagery, substantially reducing reliance on commercial satellite data. However, little SR
044 research explicitly tackles this problem or reliably addresses the accompanying distribution-shift
045 challenge. Qi et al. (2025)

046 Diffusion-based SR now dominates perceptual-quality benchmarks and offers precise structural con-
047 trol via conditional mechanisms (e.g., ControlNet Zhang & Agrawala (2023)), yet its success typ-
048 ically hinges on two assumptions that fail in scientific RS: (i) the low-resolution (LR) input is a
049 structurally faithful proxy for the desired high-resolution (HR) scene; (ii) the model’s prior has
050 seen enough semantically similar examples to render the target phenomenon. When either assump-
051 tion breaks—blurred structure or out-of-distribution (OOD) semantics—generators can hallucinate
052 textures that look plausible but violate physics and radiometry. This tension echoes the percep-
053 tion-distortion trade-off: pushing perceptual realism often degrades fidelity, and vice versa. Moser
et al. (2024); Saharia et al. (2021); Blau & Michaeli (2018)

054 In remote sensing (RS), the primary objective is scientific fidelity rather than photorealistic reconstruction Wang et al. (2022b). First, generative models must preserve precise spectral signatures, 055 a constraint that generic models routinely violate, invalidating downstream analysis. Dong et al. 056 (2021); Dou et al. (2020) Second, the ignorant prior problem becomes a critical point of failure. 057 A model cannot generate a geologically sound rare-earth feature from aesthetic principles alone; it 058 requires an implicit understanding of geomorphology. Lacking this, it produces plausible-looking 059 artifacts, not scientifically valid image. This problem is intractable for standard methods, as the very 060 rarity of the phenomena under study makes it impossible to amass the vast datasets these models 061 typically require. Chen et al. (2024); Liu et al. (2022)

063 These challenges are magnified for rare landforms, such as retrogressive thaw slumps (RTS), ice-rich 064 permafrost failures whose occurrence and expansion have intensified under Arctic warming. Their 065 rarity yields extreme class imbalance, severe few-shot regimes, and strong OOD shift across sensors 066 and regions. Any SR framework used for RTS must therefore (i) recognize OOD conditions and 067 reduce prior creativity when evidence is weak, and (ii) ground synthesis in physics, not aesthetics. 068 Nesterova et al. (2024); Lewkowicz & Way (2019); Barth et al. (2025); Nitze et al. (2025)



085 Figure 1: When the LR input is blurry or
086 semantically OOD, spatial-only guidance pre-
087 serves coarse morphology yet remains soft,
088 while semantic-only guidance hallucinates plau-
089 sible—but incorrect—textures. *RareFlow* bal-
090 ances structural evidence from the LR image with
091 textual semantics, suppressing hallucinations and
092 preserving physically consistent geometry and
093 spectra.

094
095
096
097
098
099
100
101
102
103
104
105
106
107

Confronting this triad of challenges—blurry guidance, absent priors, and data scarcity—demands more than incremental improvements; it requires a framework explicitly designed for this conflict. We therefore propose *RareFlow*, a physics-aware SR framework tailored for cross-sensor RS under severe domain shift, designed for OOD robustness. Our work reveal a fundamental tension as illustrated in Figure 1: strong spatial conditioning alone preserves coarse structure but also propagates undesirable blur, while semantic guidance alone generates realistic features but sacrifices geometric fidelity. *RareFlow* is explicitly designed to resolve this conflict. Its dual-conditioning mechanism uses a gated ControlNet (Zhang & Agrawala, 2023) to dynamically weigh the blurry source image for structural integrity while simultaneously leveraging textual prompts to provide the rich semantic context required to synthesize scientifically plausible details, even when the target phenomenon is rare and visually complex.

Contributions. The main contributions of this paper are summarized as follows.

1. **Dynamic Control of Priors to Mitigate Hallucination.** We introduce a dual-conditioning framework where a Gated ControlNet preserves geometric fidelity from the LR input, while text prompts guide semantic synthesis. This gating mechanism dynamically balances the two influences, reducing reliance on learned priors for OOD data and preventing feature hallucination.
2. **A Physics-Aware Loss for Preserving Scientific Imagery.** We design a multifaceted loss function that anchors the model’s output to the ground truth. By enforcing consistency in both the spectral domain and a perceptually uniform color space, we ensure the SR process preserves the large-scale radiometric information captured by the sensor and prevents the hallucination of erroneous detail.
3. **Unified Framework for Harmonization and SR.** We present a single, end-to-end model that jointly performs SR and radiometric harmonization. This removes the need for separate

108 pre-processing pipelines and ensures color and brightness statistics are consistent with a
 109 target reference during the reconstruction process.
 110

111 **4. State-of-the-Art (SOTA) Performance in a Low-Data Regime.** On a curated multi-
 112 sensor earth-observation benchmark emphasizing RTS and with only ≈ 800 labeled RTS
 113 images, RareFlow reconstructs fine-scale structures and surpasses strong baselines con-
 114 firmed by both quantitative metrics and qualitative analysis; an expert study further eval-
 115 uated superior fidelity and scientific integrity.

116 2 RELATED WORKS

119 Traditional SR methods, from classical interpolation to early deep learning approaches, often strug-
 120 gled to reconstruct fine, high-frequency details Johnson et al. (2016); Ledig et al. (2017); Sajjadi
 121 et al. (2017); Blau & Michaeli (2018); Moser et al. (2024). Their limitation tends to average out
 122 plausible solutions and results in overly smooth or blurry textures Blau & Michaeli (2018); Johnson
 123 et al. (2016).

124 Diffusion models address this by learning the distribution of natural images and sampling from it,
 125 enabling plausible, high-fidelity detail beyond simple sharpening Ho et al. (2020); Song et al. (2021);
 126 Saharia et al. (2023); Moser et al. (2024). This generative capability underlies SOTA perceptual SR,
 127 but applying it to remote sensing demands (i) strict geometric fidelity, (ii) scientific/radiometric plau-
 128 sibility, and (iii) robustness across sensors and styles Lanaras et al. (2018b); Gascon et al. (2017);
 129 Scarpa & Ciotola (2022a); Claverie et al. (2018); Ju et al. (2025). Our work is situated at the inter-
 130 section of these three key challenges.

131 **Enforcing Geometric and Structural Fidelity** A primary challenge in SR is maintaining strict fi-
 132 delity to the geometric and structural content of the LR input. A prominent line of work injects
 133 strong *spatial conditioning*: ControlNet-centric designs like ControlSR Zhang et al. (2023); Wan
 134 et al. (2024) enhance geometric faithfulness by injecting LR spatial cues at multiple scales. Other
 135 methods leverage segmentation masks or edge priors including SAM-DiffSR Wang et al. (2024) and
 136 SAMSRI Liu et al. (2025), to enforce sharp boundaries and object consistency. A second strategy
 137 involves using *structure-aware objectives* that directly penalize geometric distortion during train-
 138 ing. SPSR restores and supervises image gradients, using a gradient branch and gradient loss to
 139 preserve edges and linework Ma et al. (2020). To avoid hallucinations that violate the LR mea-
 140 surement, *data-consistent diffusion* methods impose the forward model during sampling: DDRM
 141 and DDNM enforce measurement consistency for SR and other linear inverse problems, while hard
 142 data-consistency with latent diffusion (ReSample) projects samples back to the measurement mani-
 143 fold at each step Kawar et al. (2022); Wang et al. (2022c); Song et al. (2024).

144 Finally, *physics-aware models* explicitly embed knowledge of the image acquisition process. These
 145 methods often incorporate the sensor’s Point Spread Function (PSF) or noise characteristics into
 146 a forward degradation model, ensuring the output is physically plausible Yang & Ren (2010). In
 147 remote sensing, this is extended by using guidance from other spectral bands or sensor physics to
 148 preserve features like building footprints Lanaras et al. (2018a); Armannsson et al. (2021); Scarpa &
 149 Ciotola (2022b). While effective for known degradations, these approaches primarily constrain the
 150 degradation process rather than the physical properties (e.g., spectral profiles) of the final HR output
 151 itself.

152 Nevertheless, a key limitation of these approaches is their reliance on robust guidance signals like
 153 accurate segmentation maps. For rare and subtly-defined geomorphological features like RTS, ac-
 154 curate segmentation masks are prohibitively expensive to create and are often brittle in practice.
 155 These features lack the clear, consistent boundaries that segmentation-heavy models rely on. This
 156 limitation motivates our use of a more flexible gated ControlNet, which provides robust structural
 157 guidance directly from the LR image without depending on fragile, external annotations.

158 **Ensuring Semantic Plausibility for Scientific Use** Beyond structural accuracy, SR outputs must be
 159 *semantically* and *scientifically* correct; plausible is not the same as correct. Recent work improves
 160 faithfulness by stabilizing the denoising trajectory and balancing LR evidence with generative pri-
 161 ors: timestep-aware fusion preserves input semantics early while enabling detail later Lin et al.
 162 (2024), and joint fine-tuning with alignment/consistency objectives reduces hallucination and better
 163 anchors content to the LR input Chen et al. (2025b). A complementary direction injects richer se-

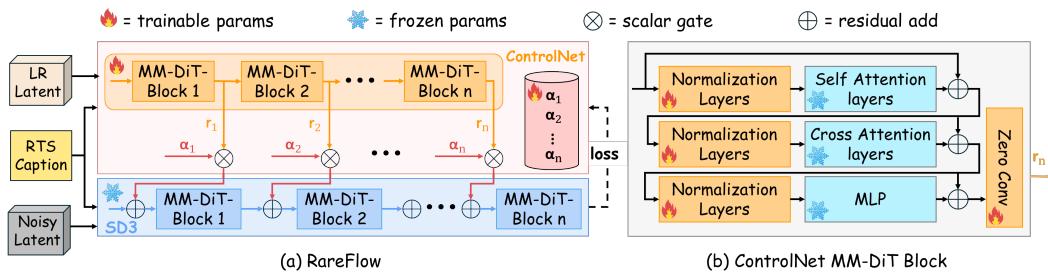


Figure 2: **(a) The training architecture of RareFlow.** The control path (orange) consumes LR latents and caption tokens to produce residual hints r_i and predicts per-block scalars $\alpha^l \in [0, 1]$ that scale r_i before injection into the frozen backbone (blue) features $F_i \leftarrow F_i + \alpha_i r_i$. **(b)** ControlNet MM-DiT block internals which produces α_i (see Sec. 3.1).

mantic priors so the model “knows what it is restoring”: cross-modal guidance via MLLMs provides scene- and object-level cues Qu et al. (2024), while segmentation-conditioned diffusion explicitly constrains regions to class-consistent appearance Xiao et al. (2024); this builds on earlier semantic SR showing that category-aware conditioning reduces implausible textures Wang et al. (2018).

However, these methods are optimized for generic photo domains. They excel at producing plausible textures for common objects but lack the domain-specific knowledge to reconstruct the unique morphology of a thaw slump or distinguish it from similar-looking terrain. Their reliance on general-world semantics risks generating photorealistic yet scientifically inaccurate artifacts. Our work directly addresses this gap by using highly specific textual prompts to encode geological knowledge, guiding the model to generate features that are not just visually convincing but also geologically plausible.

Cross-Sensor Style Transfer in SR A critical challenge in satellite imaging, often overlooked by general-purpose SR models, is the pronounced stylistic variation across sensors. Each platform exhibits distinct radiometry, spectral bandpasses, point-spread/MTF characteristics, noise statistics, and sun–sensor geometries, producing domain shifts that alter texture and color distributions Gascon et al. (2017); European Space Agency (ESA) (c); Claverie et al. (2018); Ju et al. (2025). Most multi-sensor SR and fusion methods—including recent diffusion-based pipelines such as DiffFuSR and SatDiffMoE—treat these differences as a nuisance to be harmonized or fused into a canonical representation, rather than reproducing a target sensor’s style Sarmad et al. (2025); Luo et al. (2024). In contrast, cross-sensor domain-adaptation studies in recognition consistently document substantial performance gaps attributable to sensor/style shift, underscoring that the discrepancy is persistent and non-trivial Wang et al. (2022a); Li et al. (2025); Zeng et al. (2024). Consequently, explicit *style transfer* across sensors within SR—achieving a target instrument’s textural and radiometric signature while preserving geometry—remains underexplored; doing so requires sensor-aware conditioning and physics-aware forward models (e.g., bandpass and MTF matching), together with evaluations that assess both radiometric agreement and style fidelity beyond generic distortion metrics Claverie et al. (2018); Ju et al. (2025); Scarpa & Ciotola (2022a).

In contrast, we treat sensor style as a valuable signal. We reframe the problem not as style harmonization, but as explicit style transfer within the SR task. This novel perspective allows our model to learn the specific visual signature of a target sensor (e.g., higher contrast, unique textural patterns) and intelligently apply it to an image from a different source. This reframing provides fine-grained control over the final appearance, enabling both detail enhancement and stylistic consistency with the characteristics of the chosen sensor. Our work explicitly tackles this dual objective for rare geological features, filling a crucial gap in the remote sensing literature.

3 RAREFLOW: A PHYSICS-AWARE DUAL-CONDITIONED FLOW MATCHING

We work in latent space with a frozen VAE $(\mathcal{E}, \mathcal{D})$ and a frozen *diffusion transformer* f_θ (SD3 MMDiT) Esser et al. (2024); Peebles & Xie (2023). Given an HR image y , we encode $z_0 = \mathcal{E}(y)$ and build noisy latents via the SD3 flow-match schedule

$$z_t = (1 - \sigma_t) z_0 + \sigma_t \epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}), \quad (1)$$

216 where $\{\sigma_t\}$ is the discrete schedule of the FlowMatch Euler solver. The denoiser receives (z_t, t)
 217 together with text conditioning and ControlNet residuals.
 218

219 3.1 CONTROL ADAPTERS WITH UNCERTAINTY-GATED SCALARS

221 A ControlNet g_ϕ processes the LR *latent* $\tilde{x} = \mathcal{E}(x)$ and emits per-block residual feature maps
 222 $\{r^l\}_{l \in \mathcal{L}}$ aligned with the transformer blocks Zhang & Agrawala (2023); Zavadski et al. (2024). To
 223 modulate these residuals in a *shape-preserving* way, we learn a scalar gate per block that depends
 224 on diffusion time and uncertainty:

$$225 \quad \alpha^l(t, u) = \sigma(p_0^l + p_t^l \cdot \text{norm}(t) + p_u^l \cdot u), \quad (2)$$

$$227 \quad \tilde{r}^l = s_{\text{ctrl}} \alpha^l(t, u) r^l, \quad s_{\text{ctrl}} \in \mathbb{R}_+, \quad (3)$$

229 where $\sigma(\cdot)$ is the logistic function and $\text{norm}(t) \in [0, 1]$ is a normalized time index. The scalar
 230 $u \in [0, 1]$ summarizes *only* MC-dropout uncertainty.
 231

232 3.2 UNCERTAINTY VIA MONTE CARLO DROPOUT

233 We estimate epistemic uncertainty with MC dropout Gal & Ghahramani (2016) while freezing all
 234 backbones. Dropout (rate p_{do}) is enabled only in the trainable ControlNet g_ϕ and $\alpha^l(\cdot)$. For a fixed
 235 (x, c) we draw T stochastic reconstructions
 236

$$237 \quad \hat{y}^{(k)} = \mathcal{D}\left(f_\theta(z_t, t, c, x; g_\phi \text{ with dropout})\right), \quad k = 1, \dots, T, \\ 238 \quad v(i, j) = \text{Var}_{k=1}^T [\hat{y}_{i,j}^{(k)}], \\ 239 \quad u = \text{clip}\left(\frac{\text{mean}(\text{clip}(v, 0, \tau))}{\tau}, 0, 1\right). \quad (4)$$

242 We set the scale on-the-fly as
 243

$$244 \quad \tau = \text{Pct}_{95}(\{\text{mean}(v)\}_{\text{mini-batch}}) + \epsilon, \quad (5)$$

246 or equivalently use a fixed scale $\kappa > 0$ via
 247

$$248 \quad u = 1 - \exp\left(-\frac{\text{mean}(v)}{\kappa}\right). \quad (6)$$

249 During sampling, u modulates the gates via Eq. equation 2.
 250

251 3.3 PHYSICS-AWARE TRAINING OBJECTIVE

253 We train the control pathway g_ϕ and gate parameters $\{p_0^l, p_t^l, p_u^l\}$ while keeping $(\mathcal{E}, \mathcal{D})$ and f_θ
 254 frozen. For training pairs (x, y) , define
 255

$$256 \quad \hat{y} = \mathcal{D}(f_\theta(z_t, t, c, x)), \quad (7)$$

257 and the preconditioned FlowMatch loss Lipman et al. (2022)
 258

$$259 \quad \mathcal{L}_{\text{base}} = \mathbb{E}_{(x,y), t, \epsilon} \left[\left\| \omega(\sigma_t) (f_\theta(z_t, t, c, x) - z_0) \right\|_2^2 \right], \quad (8)$$

261 with schedule-specific weight $\omega(\sigma_t)$. Although effective, this base term alone can produce overly
 262 smooth outputs; we therefore add three complementary losses.
 263

Frequency alignment (spectral magnitude). To discourage oversmoothing and limit hallucinated
 264 detail, we align spectral magnitudes in either latent or pixel space with a radial emphasis on mid/high
 265 frequencies Jiang et al. (2021); Fuoli et al. (2021):
 266

$$267 \quad \mathcal{L}_{\text{fft}} = \left\| W \odot \left(|\mathcal{F}(u_\theta)| - |\mathcal{F}(u^*)| \right) \right\|_1, \quad (9)$$

268 where u_θ and u^* denote latents, and $W(\rho) \propto \rho^\gamma$ stresses mid/high spatial frequencies that SR tends
 269 to lose.
 270

270 **Radiometric Consistency (perceptual color).** To preserve coarse color/brightness, we compare
 271 in CIELAB after blur. Let $\text{Lab}(u) \in \mathbb{R}^{H \times W \times 3}$, G_b be channelwise Gaussian blur, and $\text{Lab}_b(\cdot) \equiv$
 272 $G_b(\text{Lab}(\cdot))$. With per-channel spatial mean $\mu(\cdot)$ and stdev $\sigma(\cdot)$,

$$\mu(u) = \text{spatial mean per channel}, \quad \sigma(u) = \text{spatial stdev per channel}, \quad (10)$$

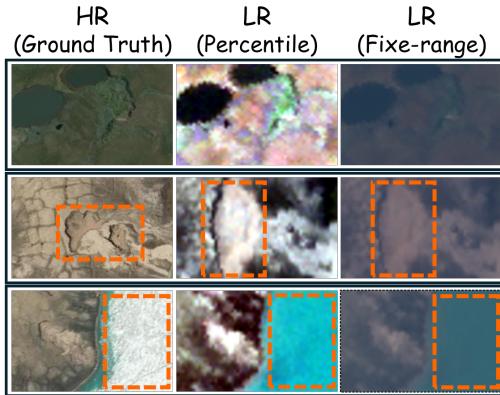
$$\mathcal{L}_{\text{color}} = \|\text{Lab}_b(\hat{y}) - \text{Lab}_b(y)\|_1 + \sum_{i \in \{\mu, \sigma\}} \|i(\text{Lab}_b(\hat{y})) - i(\text{Lab}_b(y))\|_1. \quad (11)$$

277 **Total objective.** Finally, we add LPIPS (AlexNet) Zhang et al. (2018a) to better correlate with
 278 human judgment. Our physics-aware loss enforces spectral and perceptual color consistency, pre-
 279 serving large-scale radiometry while discouraging hallucinated detail:

$$\mathcal{L} = \mathcal{L}_{\text{base}} + \lambda_{\text{fft}} \mathcal{L}_{\text{fft}} + \lambda_{\text{color}} \mathcal{L}_{\text{color}} + \lambda_{\text{lips}} \mathcal{L}_{\text{lips}}. \quad (12)$$

282 4 EXPERIMENTS AND ANALYSIS

284 We validate our approach on a challenging, real-world dataset for remote sensing SR. The task is
 285 to learn a mapping from 10m Sentinel-2 (LR) images to 2m Maxar (HR) ground truth. This cross-
 286 sensor and cross-temporal setup inherently induces OOD conditions and presents several significant
 287 challenges, illustrated in Fig. 3.



301 Figure 3: Data challenges. Left to Right: HR
 302 (Maxar), LR (Sentinel-2, Percentile Norm), LR
 303 (Sentinel-2, Fixed Norm). Row 1 shows color dis-
 304 crepancies and the effect of normalization. Row 2
 305 shows spatial misalignment. Row 3 shows tempo-
 306 ral misalignment due to changes in snow cover.

307 These challenges include: **(1) Spatio-
 308 temporal misalignment** between LR and
 309 HR images, acquired at different times,
 310 causes sub-pixel shifts, dramatic variations
 311 in illumination, and stark land cover changes.
 312 Furthermore, the dataset is characterized by
 313 **(2) Small image dimensions**, with inputs as
 314 small as 30x40 pixels, which prevent direct
 315 comparison to models evaluated on larger
 316 benchmark images. **(3) A non-standard 12-
 317 bit data range** that departs from the typ-
 318 ical 8-bit format and makes model per-
 319 formance highly sensitive to the chosen nor-
 320 malization method as it materially alters in-
 321 put data distribution and model performance.
 322 **(4) A limited training corpus of ≈ 800 im-
 323 ages**, which necessitates a data-efficient ap-
 324 proach unsuitable for training large models
 325 from scratch. Further details are provided in
 326 Appendix 1.

327 4.1 EVALUATION STRATEGY

328 We compare our method against SOTA SR models from two categories. First, we include architec-
 329 tures specifically tailored for remote sensing and trained on Sentinel-2 LR images, such as OpenSR
 330 (Donike et al., 2025), MISR-S2 (Okabayashi et al.) and ZoomLDM (Yellapragada et al., 2025).
 331 Second, we benchmark against leading general-purpose methods like AdcSr (Chen et al., 2025a),
 332 SamSR (Liu et al., 2025) and SeeSR (Wu et al., 2024) to test whether their powerful architectures
 333 offer a fundamental advantage that translates to the remote sensing domain.

334 Our evaluation strategy involves testing in two settings. We first benchmark all models on the real-
 335 world, cross-sensor pairs. We then use an idealized setting with synthetically downsampled HR
 336 images to isolate architectural performance. The quantitative metrics used for these evaluations are
 337 described and justified in detail in Appendix 2.

338 4.2 QUANTITATIVE ANALYSIS: THE FIDELITY VS. REALISM TRADE-OFF

339 **Quantitative trends.** Table 1 summarizes fidelity (PSNR/SSIM/FSIM), perceptual similarity
 340 (LPIPS/DISTS), and no-reference realism (FID/NIQE/MANIQUE). RareFlow attains the lowest

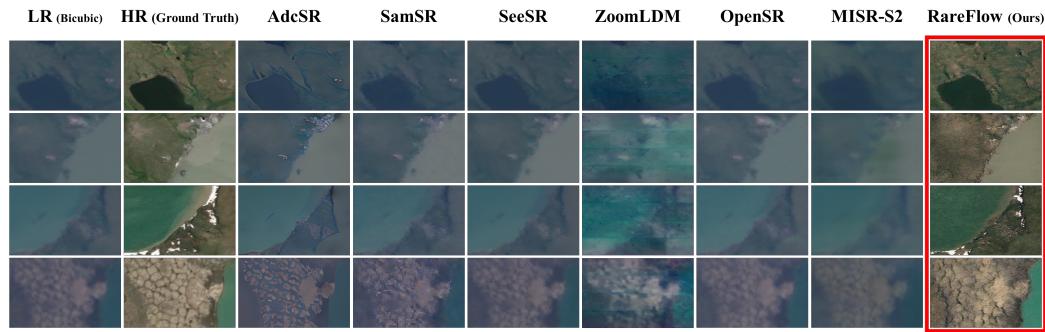


Figure 4: Qualitative comparison on paired LR–HR data.

LPIPS (0.36) and DISTS (0.30), and the lowest FID (116.16). Relative to the next-best FID (AdeSR, 187.18), this is a 38% reduction. On fidelity, RareFlow achieves the best SSIM (0.59) and FSIM (0.83) while remaining competitive in PSNR (18.76 dB vs. 18.78 dB for SeeSR.) These results are consistent with the well-known perception–distortion trade-off (Blau & Michaeli, 2018): methods optimized primarily for PSNR tend to underperform on perceptual/realism metrics, whereas RareFlow improves perceptual quality without materially degrading pixelwise fidelity.

Table 1: SR results on **paired LR–HR** data. **Bold** = best, underline = second-best; arrows indicate whether higher or lower is better.

SR Model	Fidelity Metrics			Perceptual Similarity		Realism Metrics		
	PSNR \uparrow	SSIM \uparrow	FSIM \uparrow	LPIPS \downarrow	DISTS \downarrow	FID \downarrow	NIQE \downarrow	MANIQA \uparrow
ZoomLDM	17.23	0.26	0.47	0.60	0.59	352.11	18.10	0.19
SeeSR	18.78	0.50	0.71	0.46	0.38	302.36	10.78	0.36
AdeSR	18.59	<u>0.58</u>	0.71	<u>0.40</u>	0.37	<u>187.18</u>	<u>8.38</u>	0.28
MISR-S2	18.39	0.50	0.68	0.54	0.43	254.70	13.55	<u>0.33</u>
SAMSR	18.36	0.54	<u>0.74</u>	0.48	0.39	189.01	11.84	0.32
OpenSR	17.29	0.51	0.66	0.41	<u>0.36</u>	225.62	9.80	0.25
RareFlow (Ours)	18.76	0.59	0.83	0.36	0.30	116.16	5.36	0.31

Beyond pure resolution enhancement, the core challenge of this task is performing simultaneous SR and cross-sensor style transfer. A successful model must bridge the significant domain gap between the Sentinel-2 source and the Maxar target. Figure 4 provides a striking visual demonstration of this challenge, illustrating the fundamental limitations of prior methods and the scientific utility of our approach. The baseline models fundamentally fail at the style transfer aspect. As shown in the figure, their outputs largely retain the characteristic radiometric properties of the Sentinel-2 input, such as muted contrast and a smoother textural profile. Consequently, the baselines’ attempts at SR only sharpen the wrong stylistic representation. RareFlow, however, excels at this joint task by simultaneously reconstructing geological details and translating the image into the target Maxar style, producing an output that is both HR and stylistically faithful to the ground truth.

4.2.1 CONTROLLED SR (HR \rightarrow LR DOWNSAMPLING)

To isolate the SR capacity from the cross-sensor harmonization, we synthetically downsample HR Maxar and evaluate 3-channel RGB methods under an SR protocol (details in App. .4). We exclude methods that rely on non-RGB bands, pan-sharpening inputs, auxiliary geospatial metadata, or sensor-specific priors, to ensure an apples-to-apples comparison with our RGB-only pipeline.

On HR-downsampled-HR pairs as shown in Table 2, RareFlow delivers a clearly better distortion–perception trade-off: vs. the strongest fidelity-oriented baseline (SeeSR), it cuts perceptual/realism error by –27.8% LPIPS, –21.7% DISTS, –41.7% FID, and –31.0% NIQE, while staying within 4% PSNR (SSIM ties; FSIM 1.2%). Against the closest realism challenger (SAMSR). Qualitative comparisons in the appendix .4 corroborate these trends on HR-downsampled \rightarrow HR inputs, where RareFlow reconstructs sharper textures across diverse scenes.

This experiment confirms that the core components of our architecture provide a powerful and generalizable foundation for image restoration. The leading performance in both the specialized cross-

Table 2: SR results on **HR-downsampled-HR pairs** data. **Bold** = best, underline = second-best; arrows indicate whether higher or lower is better.

SR Model	Fidelity Metrics			Perceptual Similarity		Realism Metrics		
	PSNR \uparrow	SSIM \uparrow	FSIM \uparrow	LPIPS \downarrow	DISTS \downarrow	FID \downarrow	NIQE \downarrow	MANIQA \uparrow
AdcSR	26.59	0.65	0.80	0.27	0.29	191.50	5.91	0.28
SeeSR	29.37	0.75	<u>0.86</u>	<u>0.18</u>	<u>0.23</u>	141.48	6.26	<u>0.28</u>
ZoomLDM	21.60	0.54	0.66	0.47	0.51	232.80	8.10	0.11
SAMSR	27.59	<u>0.65</u>	0.82	0.22	<u>0.23</u>	<u>128.20</u>	<u>5.20</u>	0.22
RareFlow (Ours)	<u>28.20</u>	0.75	0.87	0.13	0.18	82.53	4.32	0.46

sensor task and the standard SISR task strongly suggests that our model's design is robust and highly effective, making a significant contribution to the broader field of SR.

4.3 HUMAN EVALUATION BY DOMAIN EXPERTS

To validate the scientific utility of our results, we conducted a two-stage evaluation with geomorphology experts (details in Appendix 5).

Stage 1: Validating Semantic Guidance Integrity. Given a HR reference image with a binary mask over the region of interest, we prompt a SOTA vision–language model (GPT-5) OpenAI (2025) to produce a structured, resolution-aware description of RTS-predictive attributes (feature texture, shape, and immediate environment), see Appendix .3 for prompt design. This foundational check ensures our model is conditioned on scientifically sound semantic priors.

Stage 2: Evaluating Super-Resolved Imagery. With the integrity of the language guidance established, we proceeded to the main evaluation of our model’s SR capabilities, again leveraging domain experts. The study was designed to answer two key questions: whether our SR images offered a clarity improvement over the LR inputs, and the more challenging test of whether they could achieve perceptual parity with the HR ground truth. The results were conclusive. Experts unanimously agreed that RareFlow’s outputs provide a consistent and significant clarity enhancement over the blurry LR inputs. More critically, in numerous instances, our model achieved the gold standard: experts judged the super-resolved images to be perceptually on par with the 2m Maxar ground truth, successfully reconstructing fine-scale RTS features that were ambiguous or invisible in the source image.

4.4 ABLATION STUDIES: DISSECTING THE SOURCE OF PERFORMANCE

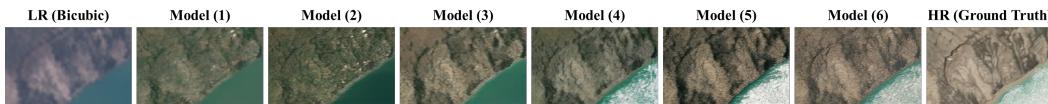
We performed comprehensive ablations (Table 3) to isolate the contributions of our key components: spatial conditioning (ControlNet), semantic guidance (text captions), and perceptual losses. Our analysis uncovers a critical dilemma inherent to this task.

We conduct a comprehensive ablation across 6 configurations to isolate the effects of (i) *spatial priors* (ControlNet: pre-trained vs. scratch, and an α -gated conditioning strength), (ii) *semantic guidance* (caption supervision describing RTS content), and (iii) *perceptual regularization* (FFT, color consistency, and LPIPS losses). The task presents unique challenges arising from two principal factors: (1) ground-truth (GT) imagery is blurred and compressed; this limitation means that classic fidelity metrics (like PSNR, SSIM) can be misleading, as a high score might only reflect a model's ability to reproduce undesirable blur, and (2) RTS scenes exhibit distinctive geomorphology. Therefore, while we report standard fidelity metrics for completeness, our analysis gives greater weight to perceptual metrics (LPIPS, DISTs) and no-reference metrics (FID, NIQE, MANIQA) which better assess the generated image's realism and visual quality.

Our analysis of Table 3 highlights a clear trade-off. The model with only a pre-trained ControlNet (2) excels at most fidelity and perceptual similarity metrics. This shows it is highly effective at replicating the GT’s structure, but it also reproduces its undesirable softness, leading to poor scores on no-reference realism metrics like MANIQA. In contrast, adding caption guidance (3, 5) boosts realism, achieving the best MANIQA scores and significantly lowering FID. However, this comes at a steep price, as fidelity metrics collapse indicating that the model is generating plausible but structurally incorrect content based on the descriptive captions.

432 Table 3: Comparison of model variants. **Bold** = best, underline = second-best; arrows indicate whether higher
 433 or lower is better.

434 435 436 437 438 439 440 441 442 443 444 445 446	434 435 436 437 438 439 440 441 442 443 444 445 446			434 435 436 437 438 439 440 441 442 443 444 445 446		434 435 436 437 438 439 440 441 442 443 444 445 446		
	434 435 436 437 438 439 440 441 442 443 444 445 446			434 435 436 437 438 439 440 441 442 443 444 445 446		434 435 436 437 438 439 440 441 442 443 444 445 446		
434 435 436 437 438 439 440 441 442 443 444 445 446								
(1) Baseline (from Scratch)	18.01	0.49	0.81	0.41	0.33	206.32	5.98	0.25
(2) + Pre-trained CN	18.80	0.51	0.85	0.35	0.29	187.60	6.42	0.19
(3) + Caption	17.08	0.38	0.74	0.47	0.32	145.43	5.64	0.32
(4) + Caption & α -gate	17.69	<u>0.52</u>	0.81	0.38	<u>0.30</u>	138.21	5.78	0.23
(5) + Caption & Loss	17.11	0.38	0.74	0.47	0.32	144.61	<u>5.61</u>	0.32
(6) Full Model (Ours)	<u>18.76</u>	0.59	<u>0.83</u>	<u>0.36</u>	<u>0.30</u>	116.16	5.36	<u>0.31</u>

447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
7710
7711
7712
7713
7714
7715
7716
7717
7718
7719
7720
7721
7722
7723
7724
7725
7726
7727
7728
7729
7730
7731
7732
7733
7734
7735
7736
7737
7738
7739
7740
7741
7742
7743
7744
7745
7746
7747
7748
7749
7750
7751
7752
7753
7754
7755
7756
7757
7758
7759
77510
77511
77512
77513
77514
77515
77516
77517
77518
77519
77520
77521
77522
77523
77524
77525
77526
77527
77528
77529
77530
77531
77532
77533
77534
77535
77536
77537
77538
77539
77540
77541
77542
77543
77544
77545
77546
77547
77548
77549
77550
77551
77552
77553
77554
77555
77556
77557
77558
77559
77560
77561
77562
77563
77564
77565
77566
77567
77568
77569
77570
77571
77572
77573
77574
77575
77576
77577
77578
77579
77580
77581
77582
77583
77584
77585
77586
77587
77588
77589
77590
77591
77592
77593
77594
77595
77596
77597
77598
77599
775100
775101
775102
775103
775104
775105
775106
775107
775108
775109
775110
775111
775112
775113
775114
775115
775116
775117
775118
775119
775120
775121
775122
775123
775124
775125
775126
775127
775128
775129
775130
775131
775132
775133
775134
775135
775136
775137
775138
775139
775140
775141
775142
775143
775144
775145
775146
775147
775148
775149
775150
775151
775152
775153
775154
775155
775156
775157
775158
775159
775160
775161
775162
775163
775164
775165
775166
775167
775168
775169
775170
775171
775172
775173
775174
775175
775176
775177
775178
775179
775180
775181
775182
775183
775184
775185
775186
775187
775188
775189
775190
775191
775192
775193
775194
775195
775196
775197
775198
775199
775200
775201
775202
775203
775204
775205
775206
775207
775208
775209
775210
775211
775212
775213
775214
775215
775216
775217
775218
775219
775220
775221
775222
775223
775224
775225
775226
775227
775228
775229
775230
775231
775232
775233
775234
775235
775236
775237
775238
775239
775240
775241
775242
775243
775244
775245
775246
775247
775248
775249
775250
775251
775252
775253
775254
775255
775256
775257
775258
775259
775260
775261
775262
775263
775264
775265
775266
775267
775268
775269
775270
775271
775272
775273
775274
775275
775276
775277
775278
775279
775280
775281
775282
775283
775284
775285
775286
775287
775288
775289
775290
775291
775292
775293
775294
775295
775296
775297
775298
775299
775300
775301
775302
775303
775304
775305
775306
775307
775308
775309
775310
775311
775312
775313
775314
775315
775316
775317
775318
775319
775320
775321
775322
775323
775324
775325
775326
775327
775328
775329
775330
775331
775332
775333
775334
775335
775336
775337
775338
775339
775340
775341
775342
775343
775344
775345
775346
775347
775348
775349
775350
775351
775352
775353
775354
775355
775356
775357
775358
775359
775360
775361
775362
775363
775364
775365
775366
775367
775368
775369
775370
775371
775372
775373
775374
775375
775376
775377
775378
775379
775380
775381
775382
775383
775384
775385
775386
775387
775388
775389
775390
775391
775392
775393
775394
775395
775396
775397
775398
775399
775400
775401
775402
775403
775404
775405
775406
775407
775408
775409
775410
775411
775412
775413
775414
775415
775416
775417
775418
775419
775420
775421
775422
775423
775424
775425
775426
775427
775428
775429
775430
775431
775432
775433
775434
775435
775436
775437
775438
775439
775440
775441
775442
775443
775444
775445
775446
775447
775448
775449
775450
775451
775452
775453
775454
775455
775456
775457
775458
775459
775460
775461
775462
775463
775464
775465
775466
775467
775468
775469
775470
775471
775472
775473
775474
775475
775476
775477
775478
775479
775480
775481
775482
775483
775484
775485
775486
775487
775488
775489
775490
775491
775492
775493
775494
775495
775496
775497
775498
775499
775500
775501
775502
775503
775504
775505
775506
775507
775508
775509
775510
775511
775512
775513
775514
775515
775516
775517
775518
775519
775520
775521
775522
775523
775524
775525
775526
775527
775528
775529
775530
775531
775532
775533
775534
775535
775536
775537
775538
775539
775540
775541
775542
775543
775544
775545
775546
775547
775548
775549
775550
775551
775552
775553
775554
775555
775556
775557
775558
775559
775560
775561
775562
775563
775564
775565
775566
775567
775568
775569
775570
775571
775572
775573
775574
775575
775576
775577
775578
775579
775580
775581
775582
775583
775584
775585
775586
775587
775588
775589
775590
775591
775592
775593
775594
775595
775596
775597
775598
775599
775600
775601
775602
775603
775604
775605
775606
775607
775608
775609
775610
775611
775612
775613
775614
775615
775616
775617
775618
775619
775620
775621
775622
775623
775624
775625
775626
775627
775628
775629
775630
775631
775632
775633
775634
775635
775636
775637
775638
775639
775640
775641
775642
775643
775644
775645
775646
775647
775648
775649
775650
775651
775652
775653
775654
775655
775656
775657
775658
775659
775660
775661
775662
775663
775664
775665
775666
775667
775668
775669
775670
775671
775672
775673
775674
775675
775676
775677
775678
775679
775680
775681
775682
775683
775684
775685
775686
775687
775688
775689
775690
775691
775692
775693
775694
775695
775696
775697
775698
775699
775700
775701
775702
775703
775704
775705
775706
775707
775708
775709
775710
775711
775712
775713
775714
775715
775716
775717
775718
775719
775720
775721
775722
775723
775724
775725
775726
775727
775728
775729
775730
775731
775732
775733
775734
775735
775736
775737
775738
775739
775740
775741
775742
775743
775744
775745
775746
775747
775748
775749
775750
775751
775752
775753
775754
775755
775756
775757
775758
775759
775760
775761
775762
775763
775764
775765
775766
775767
775768
775769
775770

486 REFERENCES
487

488 Luciano Alparone, Lucien Wald, Jocelyn Chanussot, Claire Thomas, Paolo Gamba, and Lori Mann
489 Bruce. Comparison of pansharpening algorithms: Outcome of the 2006 grs-s data-fusion contest.
490 *IEEE Transactions on Geoscience and Remote Sensing*, 45(10):3012–3021, 2007.

491 Sveinn E Armannsson, Magnus O Ulfarsson, Jakob Sigurdsson, Han V Nguyen, and Johannes R
492 Sveinsson. A comparison of optimized sentinel-2 super-resolution methods using wald’s protocol
493 and bayesian optimization. *Remote Sensing*, 13(11):2192, 2021.

494 S Barth, I Nitze, B Juhls, Alexandra Runge, and G Grosse. Rapid changes in retrogressive thaw
495 slump dynamics in the russian high arctic based on very high-resolution remote sensing. *Geo-
496 physical Research Letters*, 52(7):e2024GL113022, 2025.

497

498 Yochai Blau and Tomer Michaeli. The perception-distortion tradeoff. In *IEEE/CVF Conference on
499 Computer Vision and Pattern Recognition (CVPR)*, pp. 6228–6237, 2018. doi: 10.1109/CVPR.
500 2018.00652.

501 Bin Chen, Gehui Li, Rongyuan Wu, Xindong Zhang, Jie Chen, Jian Zhang, and Lei Zhang. Ad-
502 versarial diffusion compression for real-world image super-resolution. In *Proceedings of the
503 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 28208–28220,
504 2025a.

505

506 Junyang Chen, Jinshan Pan, and Jiangxin Dong. Faithdiff: Unleashing diffusion priors for faithful
507 image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and
508 Pattern Recognition (CVPR)*, pp. 28188–28197, 2025b.

509 Wuxing Chen, Kaixiang Yang, Zhiwen Yu, Yifan Shi, and CL Philip Chen. A survey on imbalanced
510 learning: latest research, applications and future directions. *Artificial Intelligence Review*, 57(6):
511 137, 2024.

512 Martin Claverie, Junchuan Ju, Jeffrey G. Masek, Jennifer L. Dungan, Eric F. Vermote, Jean-
513 Christophe Roger, Sergii Skakun, and Christopher Justice. The harmonized landsat and sentinel-2
514 surface reflectance data set. *Remote Sensing of Environment*, 219:145–161, 2018.

515

516 Keyan Ding, Kede Ma, Shiqi Wang, and Eero P Simoncelli. Image quality assessment: Unifying
517 structure and texture similarity. *IEEE transactions on pattern analysis and machine intelligence*,
518 44(5):2567–2581, 2020.

519 Wenqian Dong, Shaoxiong Hou, Song Xiao, Jiahui Qu, Qian Du, and Yunsong Li. Generative
520 dual-adversarial network with spectral fidelity and spatial enhancement for hyperspectral pan-
521 sharpening. *IEEE Transactions on Neural Networks and Learning Systems*, 33(12):7303–7317,
522 2021.

523

524 Simon Donike, Cesar Aybar, Luis Gómez-Chova, and Freddie Kalaitzis. Trustworthy super-
525 resolution of multispectral sentinel-2 imagery with latent diffusion. *IEEE Journal of Selected
526 Topics in Applied Earth Observations and Remote Sensing*, 2025.

527 Xinyu Dou, Chenyu Li, Qian Shi, and Mengxi Liu. Super-resolution for hyperspectral remote sens-
528 ing images based on the 3d attention-srgan network. *Remote Sensing*, 12(7):1204, 2020.

529

530 Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam
531 Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion En-
532 glish, Kyle Lacey, Alex Goodwin, Yannik Marek, and Robin Rombach. Scaling rectified flow
533 transformers for high-resolution image synthesis. *arXiv preprint arXiv:2403.03206*, 2024.

534 European Space Agency (ESA). Sentinel-2, a. URL https://www.esa.int/Applications/Observing_the_Earth/Copernicus/Sentinel-2.

535

536 European Space Agency (ESA). Worldview-3, b. URL <https://earth.esa.int/eogateway/missions/worldview-3>.

537

538 European Space Agency (ESA). S2 mission — sentiwiki (psf/mtf information). <https://sentiwiki.copernicus.eu/web/s2-mission>, c. Accessed: 2025-09-07.

539

540 Dario Fuoli, Luc Van Gool, and Radu Timofte. Fourier space losses for efficient perceptual image
 541 super-resolution. In *ICCV*, 2021.

542

543 Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model
 544 uncertainty in deep learning. In *international conference on machine learning*, pp. 1050–1059.
 545 PMLR, 2016.

546 Ferran Gascon, Catherine Bouzinac, Olivier Thépaut, Mathieu Jung, Benjamin Francesconi, Vincent
 547 Lonjou, Bruno Lafrance, Stéphane Masséra, Aurore Gaudel-Vacaresse, Frédéric Languille, Bahjat
 548 Alhammoud, Françoise Viallefont, Bringfried Pflug, Jakub Bieniarz, Stéphane Clerc, Laetitia
 549 Pessiot, Laurent Trédé, Edoardo Cadau, Roberto De Bonis, Claudio Isola, Philippe Martimort,
 550 and Victor Fernandez. Copernicus sentinel-2a calibration and products validation status. *Remote
 551 Sensing*, 9(6):584, 2017. doi: 10.3390/rs9060584.

552

553 Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter.
 554 Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in
 555 neural information processing systems*, 30, 2017.

556

557 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances
 558 in Neural Information Processing Systems (NeurIPS)*, volume 33, pp. 6840–6851, 2020.

559

560 Alain Hore and Djemel Ziou. Image quality metrics: Psnr vs. ssim. In *2010 20th international
 561 conference on pattern recognition*, pp. 2366–2369. IEEE, 2010.

562

563 Liming Jiang, Bo Dai, Wayne Wu, and Chen Change Loy. Focal frequency loss for image recon-
 564 struction and synthesis. In *ICCV*, 2021.

565

566 Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and
 567 super-resolution. In *European Conference on Computer Vision (ECCV)*, pp. 694–711. Springer,
 568 2016. doi: 10.1007/978-3-319-46475-6_43.

569

570 Junchuan Ju, Jianyang Zhuang, Zhiqiang Zhang, Grégoire Philippon, Ariuntsetseg Tsendl-Ayush,
 571 Martin Claverie, David Meyer, Crystal Schaaf, and Jeffrey G. Masek. The harmonized landsat
 572 and sentinel-2 (hls) version 2.0 surface reflectance products and subpixel co-registration strategy.
 573 *Remote Sensing of Environment*, 324:114723, 2025.

574

575 Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration
 576 models. *Advances in neural information processing systems*, 35:23593–23606, 2022.

577

578 Fred A Kruse, Adam B Lefkoff, Joseph W Boardman, Kathleen B Heidebrecht, AT Shapiro, PJ Bar-
 579 loon, and Alexander FH Goetz. The spectral image processing system (sips)—interactive visu-
 580 alization and analysis of imaging spectrometer data. *Remote sensing of environment*, 44(2-3):
 581 145–163, 1993.

582

583 Volodymyr Kuleshov, Nathan Fenner, and Stefano Ermon. Accurate uncertainties for deep learning
 584 using calibrated regression. In *International conference on machine learning*, pp. 2796–2804.
 585 PMLR, 2018.

586

587 Charis Lanaras, José Bioucas-Dias, Silvano Galliani, Emmanuel Baltsavias, and Konrad Schindler.
 588 Super-resolution of sentinel-2 images: Learning a globally applicable deep neural network. *ISPRS
 589 Journal of Photogrammetry and Remote Sensing*, 146:305–319, 2018a.

590

591 Christos Lanaras, José Bioucas-Dias, Silvano Galliani, Emmanuel Baltsavias, and Konrad Schindler.
 592 Super-resolution of sentinel-2 images: Learning a globally applicable deep neural network. *ISPRS
 593 Journal of Photogrammetry and Remote Sensing*, 146:305–319, 2018b. doi: 10.1016/j.isprsjprs.
 594 2018.09.018.

595

596 Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro
 597 Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. Photo-
 598 realistic single image super-resolution using a generative adversarial network. In *IEEE/CVF
 599 Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4681–4690, 2017. doi:
 600 10.1109/CVPR.2017.19.

594 Antoni G Lewkowicz and Robert G Way. Extremes of summer climate trigger thousands of
 595 thermokarst landslides in a high arctic environment. *Nature communications*, 10(1):1329, 2019.
 596

597 Wangbin Li, Kaimin Sun, and Jinjiang Wei. Adapting cross-sensor high-resolution remote sensing
 598 imagery for land use classification. *Remote Sensing*, 17(5):927, 2025. doi: 10.3390/rs17050927.
 599 URL <https://www.mdpi.com/2072-4292/17/5/927>.

600 Qinwei Lin, Xiaopeng Sun, Yu Gao, Yujie Zhong, Dengjie Li, Zheng Zhao, and Haoqian
 601 Wang. TASR: Timestep-aware diffusion model for image super-resolution. *arXiv preprint*
 602 *arXiv:2412.03355*, 2024.

603 Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching
 604 for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.

605 Liqin Liu, Wenyuan Li, Zhenwei Shi, and Zhengxia Zou. Physics-informed hyperspectral remote
 606 sensing image synthesis with deep conditional generative adversarial networks. *IEEE Transactions on*
 607 *Geoscience and Remote Sensing*, 60:1–15, 2022.

608 Zihang Liu, Zhenyu Zhang, and Hao Tang. Semantic-guided diffusion model for single-step image
 609 super-resolution. *arXiv preprint arXiv:2505.07071*, 2025.

610 Zhaoxu Luo, Bowen Song, and Liyue Shen. Satdiffmoe: A mixture of estimation method for satellite
 611 image super-resolution with latent diffusion models. *arXiv preprint arXiv:2406.10225*, 2024.

612 Cheng Ma, Yongming Rao, Yean Cheng, Ce Chen, Jiwen Lu, and Jie Zhou. Structure-preserving
 613 super resolution with gradient guidance. In *Proceedings of the IEEE/CVF conference on computer*
 614 *vision and pattern recognition*, pp. 7769–7778, 2020.

615 Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. Making a “completely blind” image quality
 616 analyzer. *IEEE Signal processing letters*, 20(3):209–212, 2012.

617 Brian B. Moser, Arundhati S. Shanbhag, Federico Raue, Stanislav Frolov, Sebastian Palacio, and
 618 Andreas Dengel. Diffusion models, image super-resolution and everything: A survey, 2024.

619 N. Nesterova, M. Leibman, A. Kizyakov, H. Lantuit, I. Tarasevich, I. Nitze, A. Veremeeva, and
 620 G. Grosse. Review article: Retrogressive thaw slump characteristics and terminology. *The*
 621 *Cryosphere*, 18(10):4787–4810, 2024. doi: 10.5194/tc-18-4787-2024. URL <https://tc.copernicus.org/articles/18/4787/2024/>.

622 Ingmar Nitze, Konrad Heidler, Nina Nesterova, Jonas Küpper, Emma Schütt, Tobias Hölzer, Sophia
 623 Barth, Mark J Lara, Anna K Liljedahl, and Guido Grosse. Darts: Multi-year database of ai-
 624 detected retrogressive thaw slumps in the circum-arctic permafrost region. *Scientific Data*, 12(1):
 625 1512, 2025.

626 Aimi Okabayashi, Nicolas Audebert, Simon Donike, and Charlotte Pelletier. Cross-sensor super-
 627 resolution of irregularly sampled Sentinel-2 time series. In *EARTHVISION 2024 IEEE/CVF*
 628 *CVPR Workshop. Large Scale Computer Vision for Remote Sensing Imagery*. URL <https://hal.science/hal-04552850>.

629 OpenAI. Gpt-5 api, 2025. URL <https://platform.openai.com/>. Large language model.

630 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *ICCV*, 2023.

631 Yunliang Qi, Meng Lou, Yimin Liu, Lu Li, Zhen Yang, and Wen Nie. Advancing image
 632 super-resolution techniques in remote sensing: A comprehensive survey. *arXiv preprint*
 633 *arXiv:2505.23248*, 2025.

634 Yunpeng Qu, Kun Yuan, Kai Zhao, Qizhi Xie, Jinhua Hao, Ming Sun, and Chao Zhou. XPSR: Cross-
 635 modal priors for diffusion-based image super-resolution. In *European Conference on Computer*
 636 *Vision (ECCV)*, pp. 285–303. Springer Nature Switzerland, 2024.

637 Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad
 638 Norouzi. Image super-resolution via iterative refinement. arxiv. *arXiv preprint arXiv:2104.07636*,
 639 2, 2021.

648 Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J. Fleet, and Mohammad
 649 Norouzi. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis
 650 and Machine Intelligence*, 45(4):4713–4726, 2023. doi: 10.1109/TPAMI.2022.3204461.

651

652 Mehdi S. M. Sajjadi, Bernhard Schölkopf, and Michael Hirsch. Enhancenet: Single image super-
 653 resolution through automated texture synthesis. In *IEEE International Conference on Computer
 654 Vision (ICCV)*, 2017.

655 Muhammad Sarmad, Arnt-Børre Salberg, and Michael Kampffmeyer. DiffFuSR: Super-resolution
 656 of all sentinel-2 multispectral bands using diffusion models. *arXiv preprint arXiv:2506.11764*,
 657 2025.

658

659 Giuseppe Scarpa and Matteo Ciotola. Full-resolution quality assessment for pansharpening. *Remote
 660 Sensing*, 14(8):1808, 2022a. doi: 10.3390/rs14081808.

661

662 Giuseppe Scarpa and Matteo Ciotola. Full-resolution quality assessment for pansharpening. *Remote
 663 Sensing*, 14(8):1808, 2022b.

664 Gaurav Sharma, Wencheng Wu, and Edul N Dalal. The ciede2000 color-difference formula: Im-
 665 plementation notes, supplementary test data, and mathematical observations. *Color Research
 666 & Application: Endorsed by Inter-Society Color Council, The Colour Group (Great Britain),
 667 Canadian Society for Color, Color Science Association of Japan, Dutch Society for the Study of
 668 Color, The Swedish Colour Centre Foundation, Colour Society of Australia, Centre Français de
 669 la Couleur*, 30(1):21–30, 2005.

670 Bowen Song, Soo Min Kwon, Zecheng Zhang, Xinyu Hu, Qing Qu, and Liyue Shen. Solving inverse
 671 problems with latent diffusion models via hard data consistency. In *The Twelfth International
 672 Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=j8hdRqOUhN>.

673

674 Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben
 675 Poole. Score-based generative modeling through stochastic differential equations. In *Inter-
 676 national Conference on Learning Representations (ICLR)*, 2021.

677

678 Anh D Vu, Khang V Nguyen, Bao Q Bui, and Nidal Kamel. A comprehensive survey of super-
 679 resolution remote sensing image datasets: Evolution, challenges, and future directions. *IEEE
 680 Access*, 2025.

681

682 Yuhao Wan, Peng-Tao Jiang, Qibin Hou, Hao Zhang, Jinwei Chen, Ming-Ming Cheng, and Bo Li.
 683 ControlSR: Taming diffusion models for consistent real-world image super resolution. *arXiv
 684 preprint arXiv:2410.14279*, 2024.

685

686 Chengcheng Wang, Zhiwei Hao, Yehui Tang, Jianyuan Guo, Yujie Yang, Kai Han, and Yunhe Wang.
 687 Sam-diffrsr: Structure-modulated diffusion model for image super-resolution. *arXiv preprint
 688 arXiv:2402.17133*, 2024.

689

690 Junjue Wang, Ailong Ma, Yanfei Zhong, Zhuo Zheng, and Liangpei Zhang. Cross-sensor do-
 691 main adaptation for high spatial resolution urban land-cover mapping: From airborne to space-
 692 borne imagery. *Remote Sensing of Environment*, 277:113058, 2022a. doi: 10.1016/j.rse.
 693 2022.113058. URL <https://www.sciencedirect.com/science/article/pii/S0034425722001729>.

694

695 Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Recovering realistic texture in image
 696 super-resolution by deep spatial feature transform. In *Proceedings of the IEEE conference on
 697 computer vision and pattern recognition*, pp. 606–615, 2018.

698

699 Xuan Wang, Jinglei Yi, Jian Guo, Yongchao Song, Jun Lyu, Jindong Xu, Weiqing Yan, Jindong
 700 Zhao, Qing Cai, and Haigen Min. A review of image super-resolution approaches based on deep
 701 learning and applications in remote sensing. *Remote Sensing*, 14(21):5423, 2022b.

702

703 Yinhui Wang, Jiwen Yu, and Jian Zhang. Zero-shot image restoration using denoising diffusion
 704 null-space model. *arXiv preprint arXiv:2212.00490*, 2022c.

702 Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment:
 703 from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–
 704 612, 2004.

705 Rongyuan Wu, Tao Yang, Lingchen Sun, Zhengqiang Zhang, Shuai Li, and Lei Zhang. Seesr:
 706 Towards semantics-aware real-world image super-resolution. In *Proceedings of the IEEE/CVF*
 707 *Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 25456–25467, 2024.

708 Jiahua Xiao, Jiawei Zhang, Dongqing Zou, Xiaodan Zhang, Jimmy Ren, and Xing Wei. Se-
 709 mantic segmentation prior for diffusion-based real-world super-resolution. *arXiv preprint*
 710 *arXiv:2412.02960*, 2024.

711 Lihong Yang and Jianyue Ren. Remote sensing image restoration using estimated point spread
 712 function. In *2010 International Conference on Information, Networking and Automation (ICINA)*,
 713 volume 1, pp. V1–48. IEEE, 2010.

714 Sidi Yang, Tianhe Wu, Shuwei Shi, Shanshan Lao, Yuan Gong, Mingdeng Cao, Jiahao Wang, and
 715 Yujiu Yang. Maniqa: Multi-dimension attention network for no-reference image quality assess-
 716 ment. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*,
 717 pp. 1191–1200, 2022.

718 Yili Yang, Brendan M Rogers, Greg Fiske, Jennifer Watts, Stefano Potter, Tiffany Windholz, Andrew
 719 Mullen, Ingmar Nitze, and Susan M Natali. Mapping retrogressive thaw slumps using deep neural
 720 networks. *Remote Sensing of Environment*, 288:113495, 2023.

721 Srikanth Yellapragada, Alexandros Graikos, Kostas Triaridis, Prateek Prasanna, Rajarsi Gupta, Joel
 722 Saltz, and Dimitris Samaras. Zoomldm: Latent diffusion model for multi-scale image genera-
 723 tion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*
 724 (CVPR), pp. 23453–23463, 2025.

725 Denis Zavadski, Johann-Friedrich Feiden, and Carsten Rother. Controlnet-xs: Rethinking the con-
 726 trol of text-to-image diffusion models as feedback-control systems. In *European Conference on*
 727 *Computer Vision*, pp. 343–362. Springer, 2024.

728 Junying Zeng, Yajin Gu, Chuanbo Qin, Xudong Jia, Senyao Deng, Jiahua Xu, and Huiming Tian.
 729 Unsupervised domain adaptation for remote sensing semantic segmentation with the 2d discrete
 730 wavelet transform. *Scientific Reports*, 14(23552), 2024. doi: 10.1038/s41598-024-74781-y. URL
 731 <https://www.nature.com/articles/s41598-024-74781-y>.

732 Lin Zhang, Lei Zhang, Xuanqin Mou, and David Zhang. Fsim: A feature similarity index for image
 733 quality assessment. *IEEE transactions on Image Processing*, 20(8):2378–2386, 2011.

734 Lvmin Zhang and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models.
 735 *arXiv preprint arXiv:2302.05543*, 2023.

736 Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image
 737 diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*,
 738 pp. 3836–3847, 2023.

739 Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable
 740 effectiveness of deep features as a perceptual metric. In *CVPR*, 2018a.

741 Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable
 742 effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on*
 743 *computer vision and pattern recognition*, pp. 586–595, 2018b.

750
 751 APPENDIX
 752

753 .1 DATASETS & CHALLENGES
 754

755 The dataset covers seven Arctic regions, including the Yamal and Gydan Peninsulas, Lena River, and
 Kolguev Island in Russia, along with Herschel Island, Horton Delta, Tuktoyaktuk Peninsula, and

Banks Island in Canada Yang et al. (2023). These sites provide diverse land-cover characteristics, including tundra, ice-rich permafrost bluffs, and coastal slopes, making the dataset valuable for capturing high-frequency spatial patterns in satellite imagery. The data set contains RTS annotations that were generated through careful manual digitization of Maxar imagery, supported by multi-temporal cross-checks with Sentinel-2 and ArcticDEM Yang et al. (2023).

Although the source datasets contain multiple spectral channels and auxiliary layers (e.g., NDVI, elevation), this work intentionally restricts the model’s input and output to the RGB color space. This design choice is motivated by the goal of enhancing the generalizability of the framework beyond the domain of multi-spectral remote sensing. By focusing on standard RGB data, the proposed method is more readily adaptable to other fields where training data is often limited to three channels and may be scarce, such as medical imaging or archival photograph restoration. Our method address several difficulties inherent to the real-world SR data.

To handle the 12-bit radiometric resolution of the Sentinel-2 LR images, we considered two primary normalization methods to scale pixel values to the $[0,1]$ range. Let $X \in \mathbb{R}^{H \times W \times B}$ be the input image (height H , width W , B bands), and let Y be the normalized output with the same shape.

.1.1 PER-BAND PERCENTILE NORMALIZATION

For each band $b \in \{1, \dots, B\}$, compute the lower/upper percentile values

$$\alpha_b = Q_{p_{\text{low}}}(X_{\cdot, \cdot, b}), \quad \beta_b = Q_{p_{\text{high}}}(X_{\cdot, \cdot, b}). \quad (13)$$

Then for every pixel (i, j) :

$$Y_{i,j,b} = \begin{cases} 0, & \text{if } \beta_b = \alpha_b, \\ \text{clip}\left(\frac{X_{i,j,b} - \alpha_b}{\beta_b - \alpha_b}, 0, 1\right), & \text{otherwise.} \end{cases} \quad (14)$$

Equivalently (vectorized per band):

$$Y_{\cdot, \cdot, b} = \text{clip}\left(\frac{X_{\cdot, \cdot, b} - \alpha_b}{\beta_b - \alpha_b}, 0, 1\right), \quad \text{with the same degenerate case } (\beta_b = \alpha_b). \quad (15)$$

.1.2 FIXED-RANGE NORMALIZATION

Given constants $m = \text{min_val}$ and $M = \text{max_val}$ (e.g., $m = 0, M = 3000$):

$$Y = \begin{cases} 0, & \text{if } M = m, \\ \text{clip}\left(\frac{X - m}{M - m}, 0, 1\right), & \text{otherwise,} \end{cases} \quad (16)$$

applied element-wise.

Note. $\text{clip}(z, 0, 1) = \min\{1, \max\{0, z\}\}$ is applied element-wise, and $Q_p(\cdot)$ denotes the p -th percentile computed over all pixels of the given band.

.2 EVALUATION METRICS

We evaluate SR models using a mixture of full-reference (FR), no-reference (NR), and set-level distributional metrics and we standardize implementation details to ensure fair comparison across methods. The most fundamental metric is the Peak Signal-to-Noise Ratio (PSNR) Hore & Ziou (2010). It’s simple and fast, directly measuring the pixel-wise difference between a generated image and the ground truth. While useful for gauging raw reconstruction accuracy, its major weakness is its poor correlation with human vision; an image with a high PSNR can still look unnatural or blurry to our eyes. To bridge this gap between numerical error and perceived quality, the Structural Similarity Index Measure (SSIM) Wang et al. (2004) was introduced. Instead of comparing pixels in isolation, SSIM evaluates the similarity of luminance, contrast, and structure, offering a more perceptually relevant score.

810 **PSNR** Given a reference image $x \in [0, R]^{H \times W}$ and a reconstruction y , with mean squared error
 811

$$812 \quad 813 \quad 814 \quad \text{MSE}(x, y) = \frac{1}{HW} \sum_{i=1}^H \sum_{j=1}^W (x_{ij} - y_{ij})^2, \quad (17)$$

$$815 \quad 816 \quad \text{PSNR}(x, y) = 10 \log_{10} \left(\frac{R^2}{\text{MSE}(x, y)} \right) [\text{dB}], \quad (18)$$

817 where R is the peak pixel value (e.g., $R=255$ for 8-bit).

818 **SSIM** It compares local luminance (l), contrast (c), and structure (s) between x and y :

$$820 \quad 821 \quad 822 \quad \text{SSIM}(x, y) = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)}, \quad (19)$$

823 where μ , σ^2 , and σ_{xy} are local (Gaussian-windowed) statistics; C_1 , C_2 stabilize division.

824 While these methods improved the evaluation of structural integrity, they still couldn't fully capture
 825 the complex textures and nuanced details that make an image look realistic. This limitation paved the
 826 way for learned perceptual metrics. The Learned Perceptual Image Patch Similarity (LPIPS) Zhang
 827 et al. (2018b) metric was a breakthrough, using the internal representations of a deep neural network
 828 (like VGG) to measure similarity in a way that closely mimics human judgment. Other metrics
 829 also target key visual features; FSIM (Feature Similarity Index) Zhang et al. (2011) focuses on
 830 phase congruency and gradient magnitude, which are critical to how we perceive edges and shapes.
 831 Pushing this concept further, DISTs (Deep Image Structure and Texture Similarity) Ding et al.
 832 (2020) uses a purpose-trained network to expertly balance the importance of structural correctness
 833 and textural realism, providing one of the most comprehensive full-reference evaluations available
 834 today.

835 **LPIPS** It compares deep features from a fixed backbone (e.g., VGG). For layer l with unit-
 836 normalized features $\hat{\phi}_l(\cdot)$ and learned channel weights w_l ,

$$838 \quad 839 \quad 840 \quad d_{\text{LPIPS}}(x, y) = \sum_l \frac{1}{H_l W_l} \sum_{h,w} \| w_l \odot (\hat{\phi}_l^{h,w}(x) - \hat{\phi}_l^{h,w}(y)) \|_2^2. \quad (20)$$

841 **FSIM** It weights per-pixel similarity by phase congruency (PC) and gradient magnitude (GM).
 842 With

$$843 \quad 844 \quad S_{\text{PC}} = \frac{2 \text{PC}_x \text{PC}_y + T_1}{\text{PC}_x^2 + \text{PC}_y^2 + T_1}, \quad S_{\text{G}} = \frac{2 \text{GM}_x \text{GM}_y + T_2}{\text{GM}_x^2 + \text{GM}_y^2 + T_2}, \quad (21)$$

845 the local similarity is $S_L = S_{\text{PC}}^\alpha S_{\text{G}}^\beta$, and the image-level score is

$$847 \quad 848 \quad 849 \quad \text{FSIM}(x, y) = \frac{\sum_p [\max\{\text{PC}_x(p), \text{PC}_y(p)\} S_L(p)]}{\sum_p \max\{\text{PC}_x(p), \text{PC}_y(p)\}}. \quad (22)$$

850 **DISTS** It decomposes similarity in deep space into structure and texture terms per layer l of a
 851 fixed CNN:

$$852 \quad 853 \quad 854 \quad \text{DISTS}(x, y) = \sum_l \alpha_l (1 - \rho(\phi_l(x), \phi_l(y))) + \beta_l (1 - \rho(\mu(\phi_l(x)), \mu(\phi_l(y)))), \quad (23)$$

855 where ρ is correlation and $\mu(\cdot)$ denotes channel-wise means (texture statistics).

856 A significant challenge with all these metrics is their reliance on a perfect, HR ground truth image,
 857 which is often unavailable in real-world scenarios. This necessitates the use of no-reference, or
 858 blind, quality assessors. A classic example is NIQE (Natural Image Quality Evaluator) Mittal et al.
 859 (2012), which doesn't need a reference but instead measures how an image's statistical properties
 860 deviate from those of an ideal natural image. While NIQE provides a general sense of realism, modern
 861 approaches leverage complex deep learning for more accurate blind assessments. For instance
 862 MANIQA (Multi-dimension Attention Network for IQA) Yang et al. (2022) use sophisticated Trans-
 863 former architectures to analyze images and predict a quality score that strongly aligns with human
 864 opinion.

864 **NIQE** fits a multivariate Gaussian (MVG) to natural-scene-statistics (NSS) features from pristine
 865 images, and another MVG to features from x ; quality is the (Mahalanobis-type) distance between
 866 these Gaussians:

$$868 \quad \text{NIQE}(x) = \sqrt{(\mu_n - \mu_x)^\top \left(\frac{\Sigma_n + \Sigma_x}{2} \right)^{-1} (\mu_n - \mu_x)} \quad (\text{lower is better}). \quad (24)$$

870 **MANIQA** uses a ViT backbone with transposed (channel) and Swin-based (spatial) attention
 871 blocks and a patch-weighted head to predict a scalar quality:

$$873 \quad \hat{q} = f_\theta(x). \quad (25)$$

875 Finally, evaluating a generative SR model isn't just about the quality of a single image but also about
 876 the model's ability to produce a diverse and realistic distribution of outputs. The standard for this is
 877 the Fréchet Inception Distance (FID) Heusel et al. (2017), which measures the statistical difference
 878 between the feature distributions of many generated images and real images, effectively scoring both
 879 quality and variety.

881 **(Fréchet Inception Distance (FID))** Given Inception features for a set of SR images with (μ_g, Σ_g)
 882 and reference HR images with (μ_r, Σ_r) , model each set as a Gaussian and compute the Fréchet
 883 (Wasserstein-2) distance:

$$885 \quad \text{FID} = \|\mu_r - \mu_g\|_2^2 + \text{Tr} \left(\Sigma_r + \Sigma_g - 2(\Sigma_r \Sigma_g)^{1/2} \right). \quad (26)$$

887 We include five complementary criteria that target aspects of quality not fully captured by
 888 PSNR/SSIM. **SAM** quantifies spectral fidelity by measuring the angle (in degrees) between pre-
 889 dicted and reference spectra, thus being insensitive to scale and directly penalizing spectral shape
 890 errors—crucial for multi/hyperspectral and color-constancy-sensitive tasks. ΔE_{00} (CIEDE2000)
 891 measures perceptual color difference in CIELAB with SOTA corrections for lightness, chroma, and
 892 hue interactions; it aligns with human judgments and reveals small but perceptually important color
 893 shifts that PSNR can miss. **Gaussian NLL** evaluates probabilistic predictions by scoring the en-
 894 tire predicted distribution, rewarding sharp, accurate means and penalizing both bias and misesti-
 895 mated uncertainty (σ); lower NLL means better calibrated, better-fit likelihoods. Complementing
 896 this, **ECE_{reg}** via quantile coverage assesses *calibration* of uncertainty: across target quantile lev-
 897 els, predicted quantiles should cover the empirical outcomes at the stated rates; deviations indicate
 898 over/under-confidence even when point accuracy is high. Finally, for pansharpening and related fu-
 899 sion, **QNR** jointly measures spectral consistency across bands (D_λ) and spatial detail preservation
 900 relative to the panchromatic guide (D_s), yielding a no-reference score that balances “do no spectral
 901 harm” with “add the right spatial detail.” Alparone et al. (2007); Kuleshov et al. (2018); Sharma
 902 et al. (2005)

903 **Spectral Angle Mapper (SAM, in degrees).** For two spectra $\mathbf{x}, \mathbf{y} \in \mathbb{R}^B$,

$$905 \quad \text{SAM}(\mathbf{x}, \mathbf{y}) = \frac{180}{\pi} \arccos \left(\frac{\mathbf{x}^\top \mathbf{y}}{\|\mathbf{x}\|_2 \|\mathbf{y}\|_2} \right). \quad (27)$$

907 Kruse et al. (1993)

909 **ΔE_{00} (CIEDE2000 color difference).** Given two CIELAB colors (L_1^*, a_1^*, b_1^*) and (L_2^*, a_2^*, b_2^*) ,
 910 define chroma $C_i^* = \sqrt{(a_i^*)^2 + (b_i^*)^2}$, mean $\bar{C}^* = \frac{1}{2}(C_1^* + C_2^*)$, and
 911

$$914 \quad G = \frac{1}{2} \left(1 - \sqrt{\frac{(\bar{C}^*)^7}{(\bar{C}^*)^7 + 25^7}} \right), \quad a'_i = (1 + G)a_i^*, \quad C'_i = \sqrt{a'^2 + b'^2}, \quad h'_i = \text{atan2}(b_i^*, a_i'). \quad (28)$$

917 Let $\Delta L' = L_2^* - L_1^*$, $\Delta C' = C'_2 - C'_1$, and

918
919
920
921
922
923
924
925
926
927
928
929
930

$$\Delta h' = \begin{cases} h'_2 - h'_1 & \text{if } |h'_2 - h'_1| \leq 180^\circ, \\ h'_2 - h'_1 - 360^\circ & \text{if } h'_2 - h'_1 > 180^\circ, \\ h'_2 - h'_1 + 360^\circ & \text{if } h'_2 - h'_1 < -180^\circ, \end{cases} \quad \Delta H' = 2\sqrt{C'_1 C'_2} \sin\left(\frac{\Delta h'}{2}\right). \quad (29)$$

With means $\bar{L}' = \frac{1}{2}(L'_1 + L'_2)$, $\bar{C}' = \frac{1}{2}(C'_1 + C'_2)$, and

$$\bar{h}' = \begin{cases} \frac{h'_1 + h'_2}{2} & \text{if } |h'_1 - h'_2| \leq 180^\circ, \\ \frac{h'_1 + h'_2 + 360^\circ}{2} & \text{if } |h'_1 - h'_2| > 180^\circ \text{ and } h'_1 + h'_2 < 360^\circ, \\ \frac{h'_1 + h'_2 - 360^\circ}{2} & \text{otherwise,} \end{cases} \quad (30)$$

the weighting functions are

$$S_L = 1 + \frac{0.015(\bar{L}' - 50)^2}{\sqrt{20 + (\bar{L}' - 50)^2}}, \quad S_C = 1 + 0.045\bar{C}', \quad S_H = 1 + 0.015\bar{C}'T, \quad (31)$$

$$T = 1 - 0.17 \cos(\bar{h}' - 30^\circ) + 0.24 \cos(2\bar{h}') + 0.32 \cos(3\bar{h}' + 6^\circ) - 0.20 \cos(4\bar{h}' - 63^\circ), \quad (32)$$

$$R_C = 2\sqrt{\frac{(\bar{C}')^7}{(\bar{C}')^7 + 25^7}}, \quad R_T = -R_C \sin(2\Delta\theta), \quad \Delta\theta = 30^\circ \exp\left[-\left(\frac{\bar{h}' - 275^\circ}{25^\circ}\right)^2\right]. \quad (33)$$

Finally, for parametric factors $k_L = k_C = k_H = 1$ (unless otherwise stated),

$$\Delta E_{00} = \sqrt{\left(\frac{\Delta L'}{k_L S_L}\right)^2 + \left(\frac{\Delta C'}{k_C S_C}\right)^2 + \left(\frac{\Delta H'}{k_H S_H}\right)^2 + R_T \frac{\Delta C'}{k_C S_C} \frac{\Delta H'}{k_H S_H}}. \quad (34)$$

Sharma et al. (2005)

Gaussian Negative Log-Likelihood (per sample). For targets y_i and Gaussian predictions $\mathcal{N}(\mu_i, \sigma_i^2)$,

$$\text{NLL} = \frac{1}{2} \sum_{i=1}^n \left[\log(2\pi\sigma_i^2) + \frac{(y_i - \mu_i)^2}{\sigma_i^2} \right]. \quad (35)$$

Expected Calibration Error (ECE) for regression via quantile coverage. Let $\{\alpha_m\}_{m=1}^M \subset (0, 1)$ be nominal quantile levels and $\hat{q}_{\alpha_m}(x)$ the model's predicted α_m -quantile for input x . Define the empirical coverage at level α_m by

$$\widehat{\text{cov}}(\alpha_m) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{y_i \leq \hat{q}_{\alpha_m}(x_i)\}. \quad (36)$$

Then an ECE-style scalar summary is

$$\text{ECE}_{\text{reg}} = \sum_{m=1}^M w_m |\widehat{\text{cov}}(\alpha_m) - \alpha_m|, \quad w_m \geq 0, \quad \sum_{m=1}^M w_m = 1, \quad (37)$$

(e.g., $w_m = \frac{1}{M}$). Lower is better; 0 indicates calibrated quantiles. Kuleshov et al. (2018)

972 **QNR (Quality with No Reference) for pansharpening, and its components.** Given the fused
 973 multispectral image $\hat{\mathbf{X}} = \{\hat{X}_b\}_{b=1}^B$, the original multispectral image $\mathbf{X} = \{X_b\}_{b=1}^B$ (upsampled to
 974 the fused resolution), and the panchromatic image P , define Wang–Bovik’s universal image quality
 975 index $Q(\cdot, \cdot)$ applied bandwise (and to gradient images for spatial terms). Spectral distortion:

$$977 \quad D_\lambda = \frac{2}{B(B-1)} \sum_{1 \leq i < j \leq B} |Q(X_i, X_j) - Q(\hat{X}_i, \hat{X}_j)|. \quad (38)$$

985 Spatial distortion:

$$990 \quad D_s = \frac{1}{B} \sum_{b=1}^B |Q(\nabla X_b, \nabla P) - Q(\nabla \hat{X}_b, \nabla P)|. \quad (39)$$

995 With exponents $\alpha, \beta > 0$ (often $\alpha = \beta = 1$),

$$1000 \quad \text{QNR} = (1 - D_\lambda)^\alpha (1 - D_s)^\beta, \quad \text{higher is better.} \quad (40)$$

1004 Alparone et al. (2007)

1008 **Implementation references** We rely on widely used implementations: LPIPS (official PyTorch),
 1009 DISTS (official), MANIQA (official).¹

1012 .3 PROMPTING DETAILS FOR CAPTION GENERATION

1015 This section details the complete system prompt provided to the Vision-Language Model (GPT-
 1016 5) (OpenAI, 2025) to generate descriptive captions from the HR reference images. The objective
 1017 was to create semantically rich, resolution-aware descriptions of RTS while adhering to a natural,
 1018 non-technical style suitable for guiding the SR process.

1019 The prompt established an expert persona for the VLM and outlined a series of rules and constraints
 1020 organized into style, content, and exclusions.

1024 ¹[https://github.com/richzhang/PerceptualSimilarity \(LPIPS\)](https://github.com/richzhang/PerceptualSimilarity), [https://github.com/dingkeyan93/DISTS \(DISTs\)](https://github.com/dingkeyan93/DISTS), [https://github.com/IIGROUP/MANIQA \(MANIQA\)](https://github.com/IIGROUP/MANIQA).

1026
1027**System Prompt Fed to the VLM**

1028

Persona & Objective

1029

You are an expert satellite image analyst writing captions. Your task is to describe images of permafrost thaw slumps. The goal is to create natural, descriptive captions suitable for training a text-to-image AI model. The captions must sound like a human describing a photo in simple terms.

1030

1031

1032

1033

1034

Style Rules

1035

1036

1037

1038

1039

1040

1041

- Use simple, everyday language.
- Write in a natural, fluid style.
- Use the present tense.
- Do not refer to "this image" or "the photo".

Content Requirements

1042

1043

1044

1045

1046

1047

1048

1049

1050

1051

1052

1053

- **Main Feature:** Describe the thaw slump using common terms like "landslide," "thaw slump," "ground collapse," or "erosion scar."
- **Shape & Form:** Mention its shape with simple descriptions like "crescent-shaped," "bowl-shaped," or "tongue of dirt."
- **Colors & Textures:** Describe the colors and textures of the ground, vegetation, and water (e.g., "dark brown soil," "green tundra," "cracked earth," "blue-green ocean").
- **Setting:** Briefly describe the surrounding environment, such as "coastal cliff," "green hillside," "tundra plain," or "riverbank."

Exclusion Criteria (Crucial Constraints)

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079



Figure 6: Example of input image given to VLM to generate an RTS-aware caption.

An example of a caption generated for the given input image 6:

"A bowl-shaped thaw slump cuts into a rugged coastal hillside, exposing dark brown soil and crumbly earth, with pale sandy streaks sliding downslope toward blue-green water, surrounded by gray-brown, sparsely vegetated tundra."

.4 VALIDATION ON STANDARD SR

To rigorously assess the architectural robustness and general applicability of our proposed model, we conducted an additional set of experiments outside the primary cross-sensor domain. The objective of this analysis was to determine if the strong performance of our model is specific to the cross-sensor challenge or if its underlying design principles are fundamentally effective for the general task of Single Image Super-Resolution (SISR).

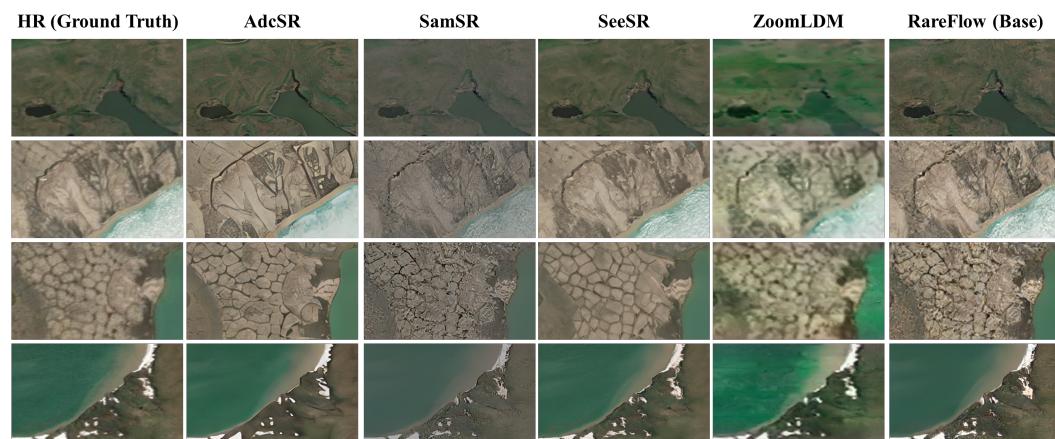


Figure 7: Qualitative comparison on HR-downsampled to HR data.

1096 .4.1 EXPERIMENTAL SETUP

1098 We used the same HR ground truth images used in our main evaluation. For this SISR task, the cor-
 1099 responding LR inputs were synthesized by bicubically downsampling the HR images by a factor of
 1100 4x. The visual results in Figure 7 corroborate the quantitative findings. **RareFlow** excels at recon-
 1101 structing fine-grained textures and sharp, plausible details that are often lost or blurred by competing
 1102 methods. While other models tend to produce overly smoothed results, our model generates clean,
 1103 realistic, and highly detailed images that are perceptually more convincing.

1104 .5 DOMAIN EXPERT EVALUATION

1106 The primary objective of this evaluation was to rigorously assess our method across three key cri-
 1107 teria: (1) the perceptual parity of our SR images with HR ground truth, (2) the clarity enhancement
 1108 over LR inputs, and (3) the quality of the VLM-generated semantic guidance.

1109 The evaluation was performed by a panel of domain experts to ensure the findings are grounded
 1110 in practical scientific application. All participants are scientists actively conducting research on
 1111 permafrost geomorphology and Arctic remote sensing at well-known international research centers.
 1112 The cohort was composed of:

- 1114 • **Research Scientists**, all holding a Ph.D. in a relevant field.
- 1115 • **Senior Research Assistants**, with extensive, specialized experience in analyzing thaw
 1116 slump features from satellite imagery.

1118 The evaluation was conducted via a custom web-based interface (Figure 8). For each of the **30 sam-**
 1119 **ples** assigned to a reviewer, the interface displayed a comprehensive view containing four images:
 1120 the HR ground truth with an RTS mask, the unmasked HR ground truth, our generated SR output,
 1121 and the original Sentinel-2 LR input. The VLM-generated caption was displayed prominently along-
 1122 side. This setup, while time-intensive for the experts, allowed for a thorough and direct comparison
 1123 of all relevant data.

1124 .6 USE OF LARGE LANGUAGE MODELS

1126 Large language models were utilized for grammatical correction, LaTeX formatting, debugging, and
 1127 finding related work.

1134

1135

1136

1137

1138

1139

1140

1141

1142

1143

1144

1145

1146

1147

1148

1149

1150

1151

1152

1153

1154

1155

1156

1157

1158

1159

1160

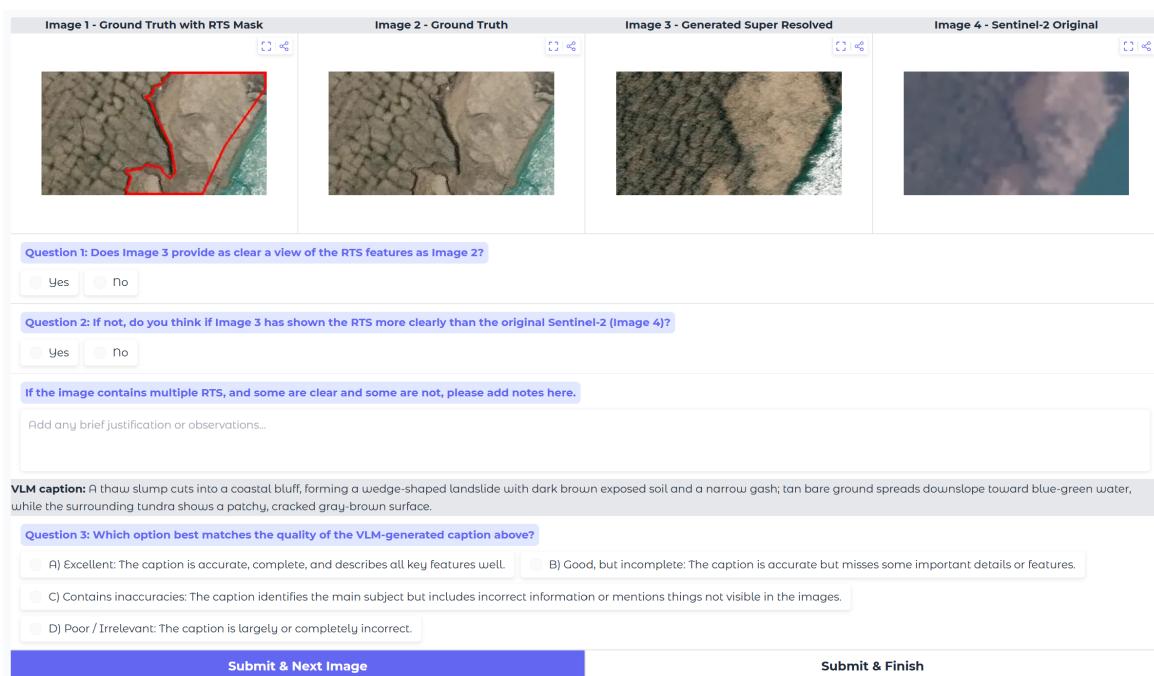
1161

1162

1163

1164

1165



1170

1171

Figure 8: Our custom web-based interface for human evaluation.

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187