Active Learning for Conditional Generative Compressed Sensing

Alexander DeLise[1,2], Nick Dexter[1]

[1] Department of Scientific Computing, Florida State University, Tallahassee, FL [2] Department of Mathematics, Florida State University, Tallahassee, FL

Problem Setup

Let $f \in \mathbb{R}^{3N}$ be a vectorized $n \times n$ RGB image, where $N = n^2$. Let

$$G : \mathbb{R}^k \times \mathbb{R}^\ell \to \mathbb{R}^{3N}, \qquad (z, \tau(c)) \mapsto f,$$

be a **conditional generative model**, where $z \in \mathbb{R}^k$ denotes a latent variable and $\tau : \mathcal{C} \to \mathbb{R}^\ell$ is a text encoder mapping a prompt $c \in \mathcal{C}$ to a conditioning embedding $\tau(c)$.

f

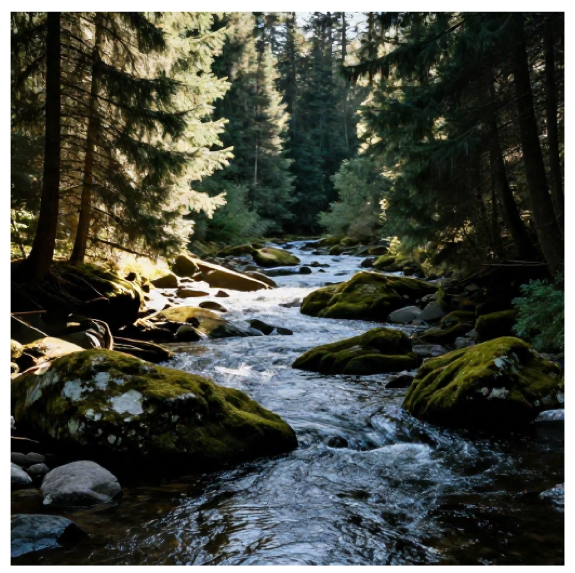

Generation Prompt

c = "Forest river bend, clear water over stones, mossy rocks, sunbeams through evergreens, realistic foliage. Pro photo, 35mm, f/5.6, ISO 200, natural contrast, high detail."

Let $F \in \mathbb{C}^{N \times N}$ be the discrete Fourier transform (DFT). Select $m \ll N$ frequencies $\Omega \subset D = \{1, \ldots, N\}$ using $P_\Omega \in \mathbb{R}^{m \times N}$.

$A^\dagger y$

Measurement

$$A = \tfrac{1}{\sqrt{m}}(I_3 \otimes P_\Omega F)$$
$$y = Af + \varepsilon$$
$$\varepsilon \sim \mathcal{N}(0, \sigma_y^2 I)$$
$$R = m/N = 0.01$$

Central Reconstruction Problem

For fixed prompt $\hat{c}$ and measurement y, we solve

$$\hat{z} = \arg\min_z \underbrace{\tfrac{1}{2\sigma_y^2}\left\| y - A\,G\big(z, \tau(\hat{c})\big) \right\|_2^2}_{\text{data consistency (DC)}} + \underbrace{\lambda \mathcal{R}_\theta(z; \hat{c})}_{\text{realism}}$$

to recover $\hat{f} = G\big(\hat{z}, \tau(\hat{c})\big)$. The DC term matches measurements to y, while the realism term (e.g. $-\log p_\theta(z \mid \hat{c})$) enforces a natural, prompt-consistent image from G.

Christoffel Sampling

Generalized Christoffel Function $K_{\mathbb{F}}$

Let $\mathbb{F} = \mathrm{Range}(G(\cdot, c))$ be our generator-induced *model class*. For Fourier index $i \in D$,

$$K_{\mathbb{F}}(i) = \sup_{f \in \mathbb{F} \setminus \{0\}} \left\| (I_3 \otimes e_i^\top F) f \right\|_2^2 \,/\, \|f\|_2^2,$$

which measures how strongly images in $\mathbb{F}$ activate frequency i.

Christoffel Sampling with $K_{\mathbb{F}-\mathbb{F}}$

We leverage CS4ML [1] to construct a sampling measure defined on the *difference class* $\mathbb{F} - \mathbb{F} = \{f_1 - f_2 : f_1, f_2 \in \mathbb{F}\}$, capturing pairwise variability in the generator range.

Sampling Rule

Using an approximation $\widetilde{K} \approx K_{\mathbb{F}-\mathbb{F}}$ from [1, Alg. 1], we build the **optimal sampling measure** $\mu_i = \widetilde{K}(i) \,/\, \sum_{j=1}^N \widetilde{K}(j)$, and sample $\Omega \sim \mu$.

Latent Diffusion Models

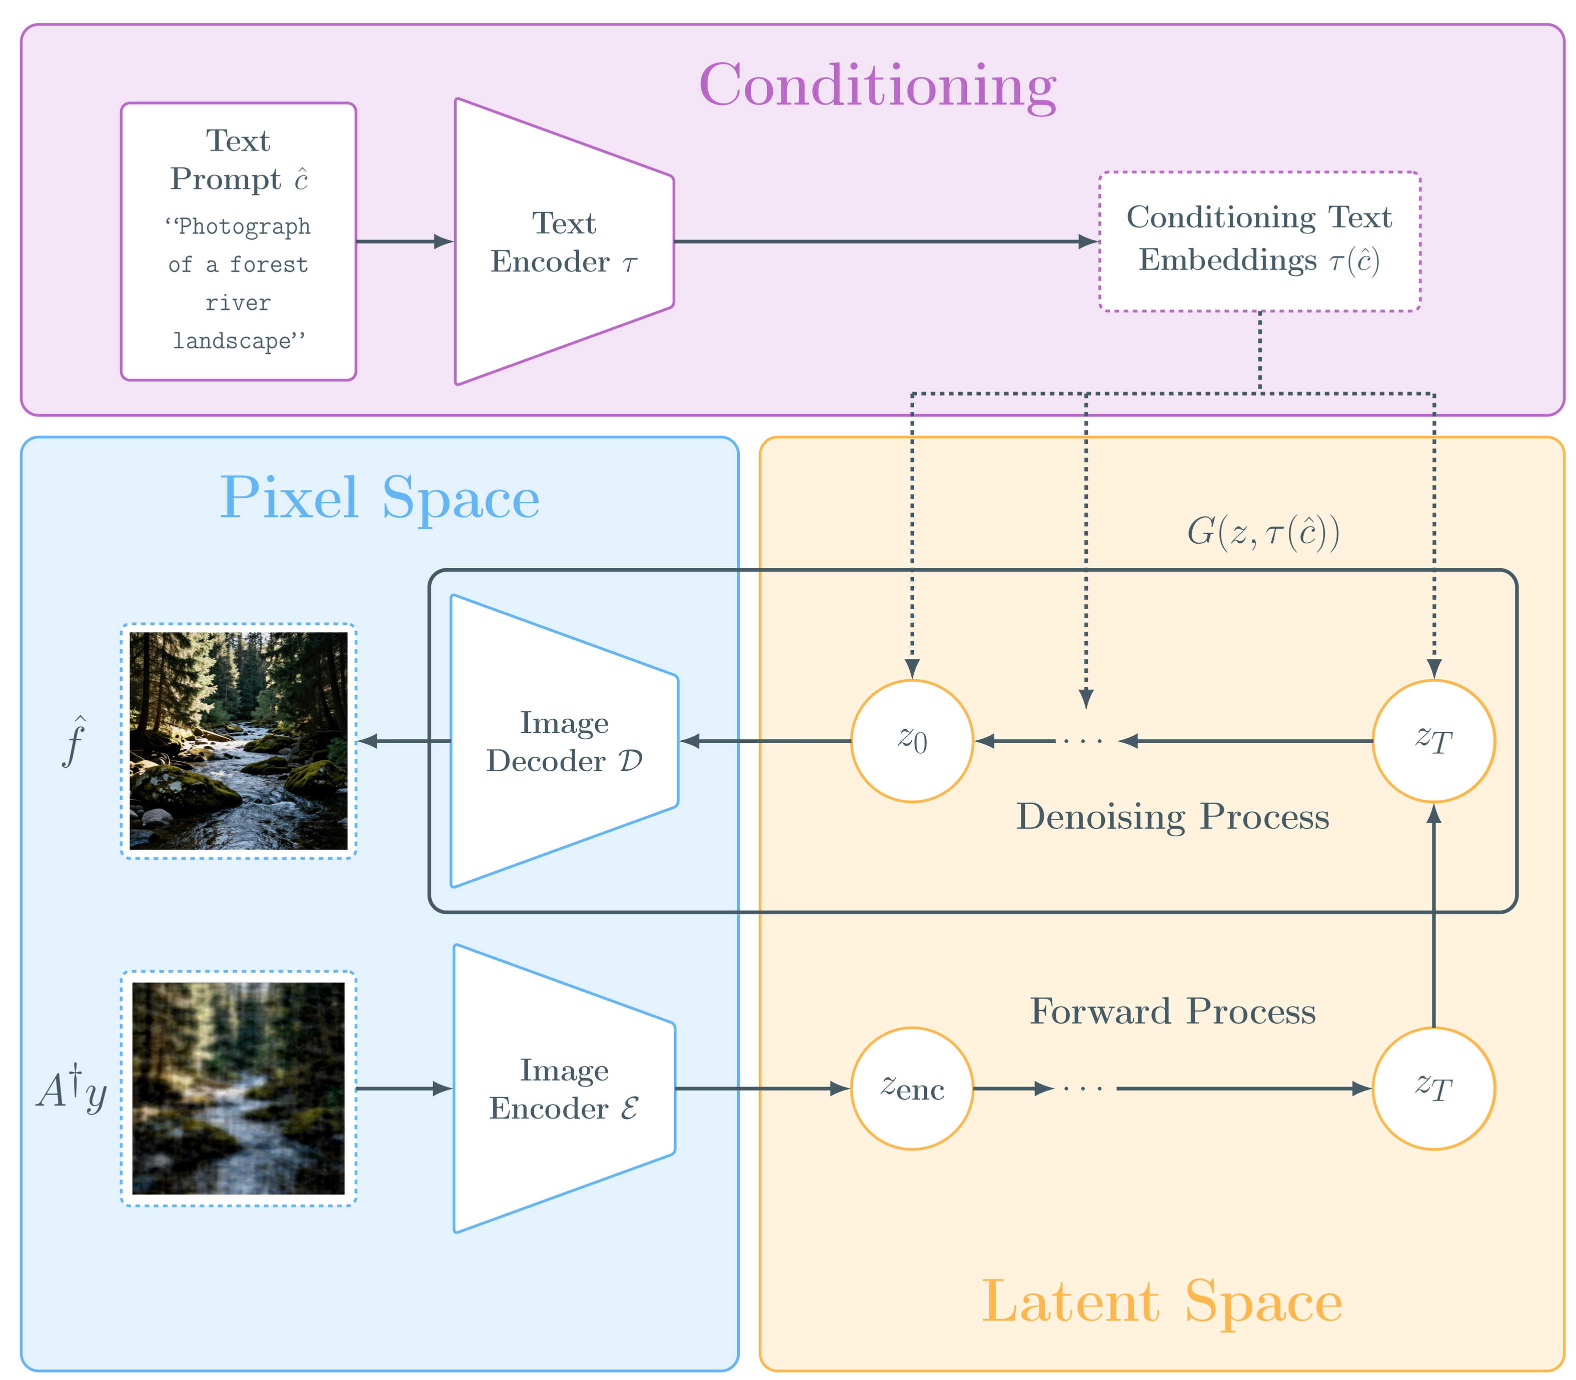

Algorithm 1: Image Reconstruction

1. Initialize a noisy latent z by encoding $A^\dagger y$.
2. Perform denoising step(s).
3. Decode and enforce DC via DDNM projection [2].
4. Encode and re-noise for Time Travel (TT).
5. Repeat steps 2–4 to reconstruct $\hat{f}$.

Sampling and Conditioning

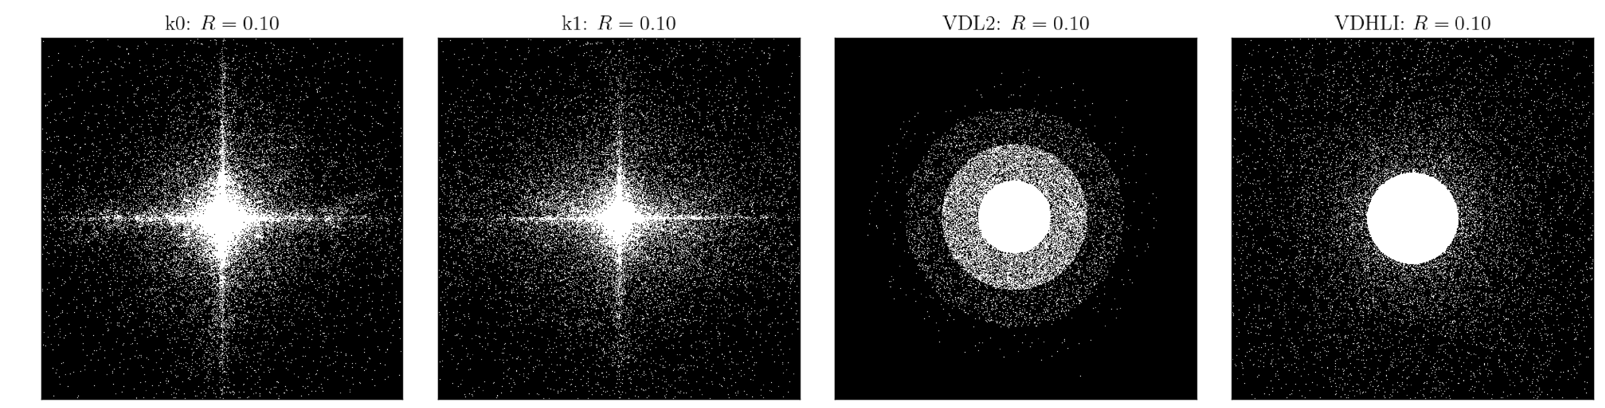

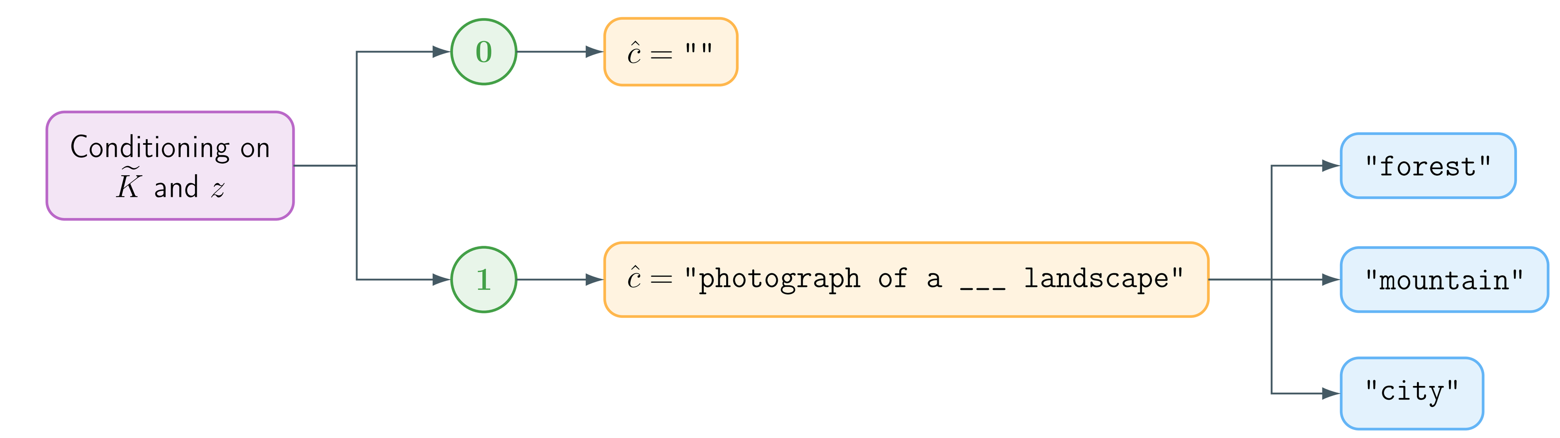

Results

PSNR

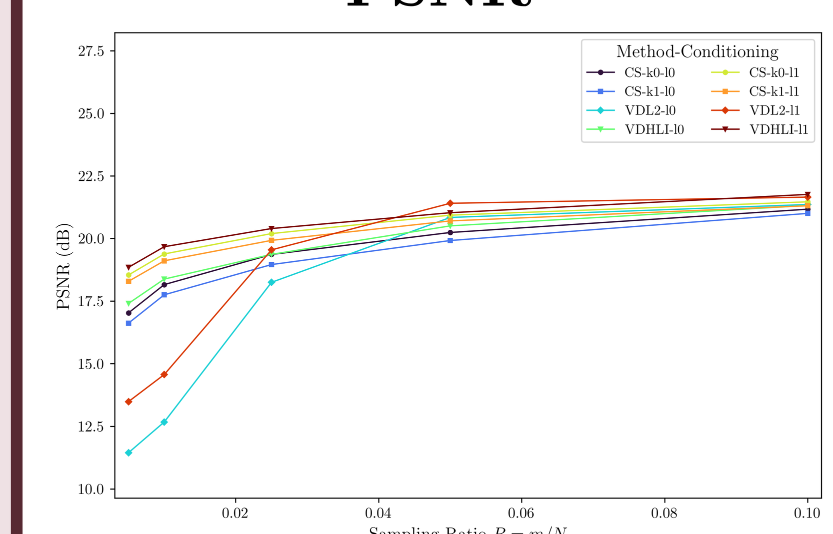

SSIM

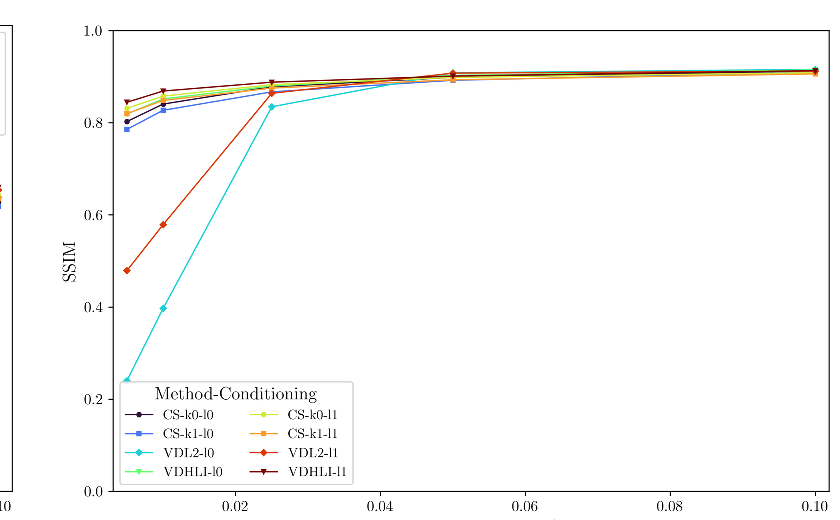

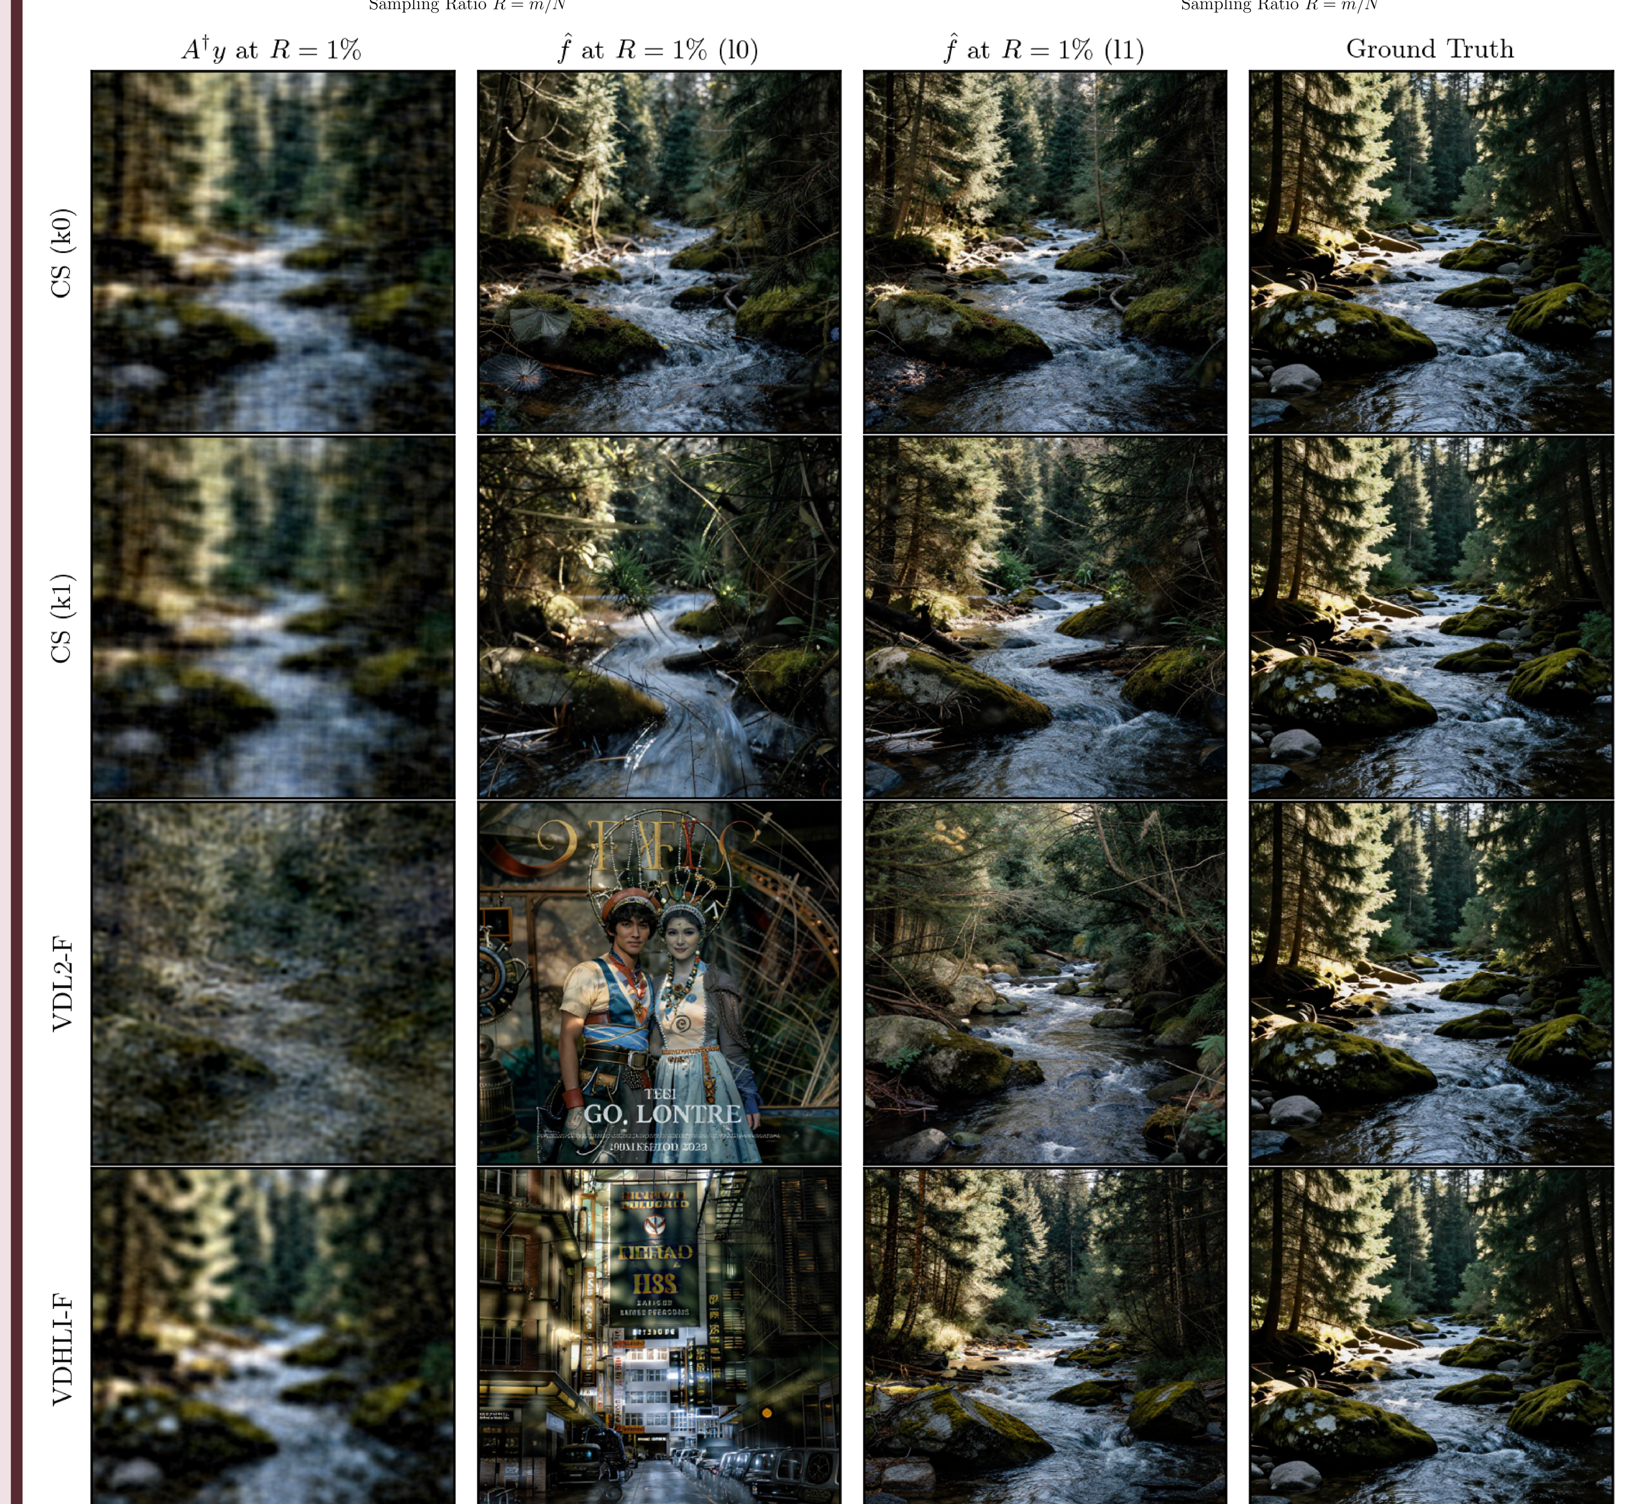

References

[1] Ben Adcock, Juan M. Cardenas, and Nick Dexter. CS4ML: a general framework for active learning with arbitrary data based on Christoffel functions. In *Advances in Neural Information Processing Systems*, 2023.

[2] Yinhuai Wang, Jiwen Yu, and Jian Zhang. Zero-shot image restoration using denoising diffusion null-space model. In *International Conference on Learning Representations*, 2023.