

# gQIR: Generative Quanta Image Reconstruction

Aryan Garg  
University of Wisconsin-Madison  
agarg54@wisc.edu

Sizhuo Ma  
Snap Inc.  
sma@snapchat.com

Mohit Gupta  
University of Wisconsin-Madison  
mgupta37@wisc.edu

## Abstract

*Capturing high-quality images from only a few detected photons is a fundamental challenge in computational imaging. Single-photon avalanche diode (SPAD) sensors promise high-quality imaging in regimes where conventional cameras fail, but raw quanta frames contain only sparse, noisy, binary photon detections. Recovering a coherent image from a burst of such frames requires handling alignment, denoising, and demosaicing (for color) under noise statistics far outside those assumed by standard restoration pipelines or modern generative models. We present an approach that adapts large text-to-image latent diffusion models to the photon-limited domain of quanta burst imaging. Our method leverages the structural and semantic priors of internet-scale diffusion models while introducing mechanisms to handle Bernoulli photon statistics. By integrating latent-space restoration with burst-level spatio-temporal reasoning, our approach produces reconstructions that are both photometrically faithful and perceptually pleasing, even under high-speed motion. We evaluate the method on synthetic benchmarks and new real-world datasets, including the first color SPAD burst dataset and a challenging Deforming (XD) video benchmark. Across all settings, the approach substantially improves perceptual quality over classical and modern learning-based baselines, demonstrating the promise of adapting large generative priors to extreme photon-limited sensing. Website: <https://aryan-garg.github.io/gqir>.*

## 1. Introduction

Capturing a clear image from just a few detected photons is a long-standing challenge in computational imaging. Single-photon avalanche diode (SPAD) sensors [40, 42] hold the promise of high-fidelity imaging in extreme low-light and high-speed regimes, where conventional sensors fail. However, each photon detection is a discrete stochastic event; as a result, individual *quanta frames* are dominated by shot noise and quantization artifacts, often containing only sparse scene information [15, 78]. Recovering a high-quality photograph from such frames requires combining information across a temporal burst of quanta frames.

SPAD arrays can operate in a Bernoulli mode, where

each pixel records a binary value: 1 if one or more photons are detected during an exposure, and 0 otherwise. A short sequence of such binary frames—termed a *nano-burst*—can be aggregated into a higher bit-depth representation (e.g., seven binary frames combined into a 3-bit frame). This burst-mode strategy enables photon-limited imaging at high frame rates and provides an opportunity to recover high-quality images from sparse photon events. However, aggregating these frames into a coherent image is challenging: small inter-frame motions cause misalignment, while the paucity of photons renders conventional motion estimation unreliable. This interplay between motion estimation, alignment, and photon-limited denoising defines the core challenge of *quanta burst reconstruction*.

Early offline approaches to this problem relied on classical vision techniques [37, 78], which explicitly estimate motion between frames. More recently, learning-based methods have been introduced [8, 9], which treat alignment and fusion as learnable modules. Despite this progress, photon-starved scenes with extreme deformation or ultra-high-speed motion remain challenging, as illustrated in Fig. 1. Leveraging the representational power of large-scale generative models for quanta imaging remains an open direction. In particular, learning-based methods do not yet exploit the structural knowledge embedded in large text-to-image (T2I) diffusion models [6, 12, 14, 32, 44, 45, 47, 48, 71]. Generative restoration models [19, 33, 34, 62, 64, 70, 72] leveraging such T2I priors have shown strong performance on conventional camera images. However, these models break down in the photon-limited regime for quanta cameras, where noise is non-Gaussian and photon counts far below those in conventional photography. Naive fine-tuning of these models leads to shortcut learning and does not produce meaningful outputs. Furthermore, most prior quanta reconstruction methods consider monochrome sensors. In contrast, photon-counting color sensors [37] introduce additional challenges due to sparse photon events in each color channel. Together, these factors make photon-limited burst reconstruction a demanding testbed for adapting large generative models to discrete, sparse quanta measurements. As shown in Fig. 1, the benefits of doing so emerge clearly under extreme de-

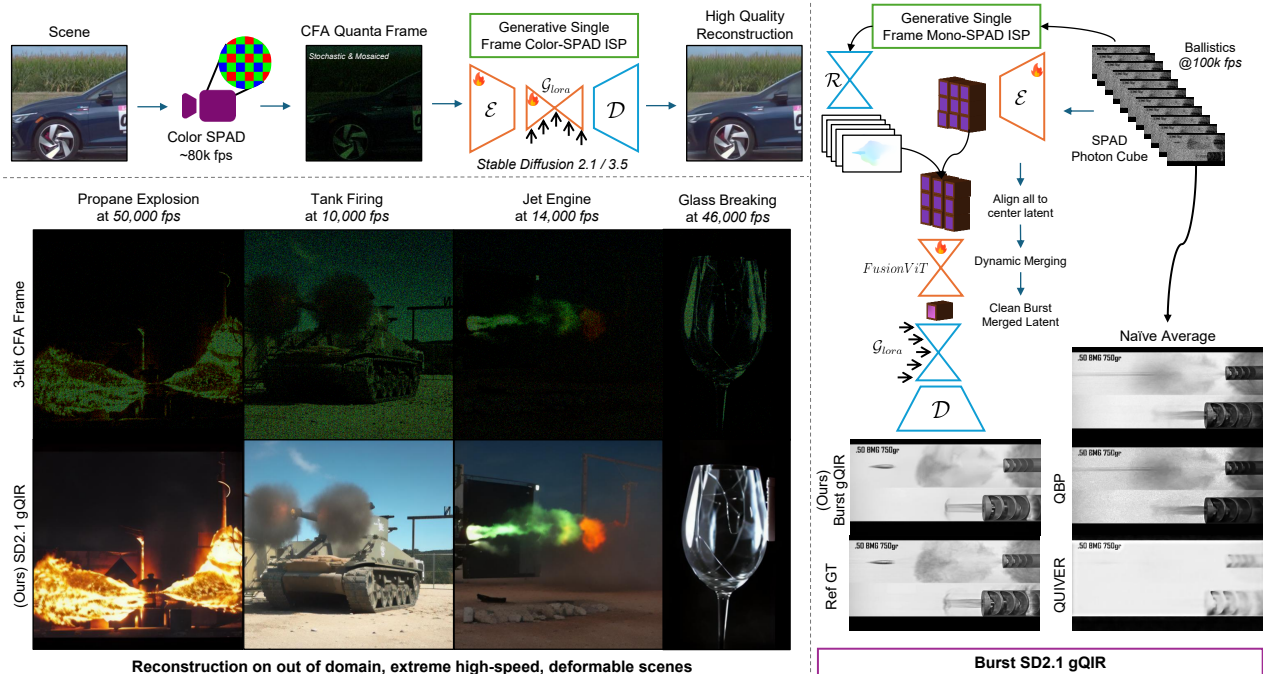


Figure 1. **gQIR: Photorealistic single image and burst reconstruction from ultra-high-speed color SPADs.** Our pipeline reconstructs high-quality RGB images from 3-bit color-SPAD CFA nano-bursts (left) and merges SPAD photon cubes into temporally consistent bursts (right). From photon-starved inputs captured at 10k–50k fps in extreme, out-of-domain scenes, gQIR recovers sharp textures, accurate color, and coherent structure by leveraging a generative prior. For burst sequences up to 100k fps, FusionViT aligns and dynamically merges quanta latents, outperforming traditional and learning-based methods in both fidelity and perceptualness under motion.

formation or ultra-high-speed conditions.

To address these challenges, we propose a modular three-stage framework that adapts latent diffusion models for quanta burst reconstruction. The first stage jointly denoises and demosaics single or nano-burst quanta frames by finetuning the variational autoencoder (VAE) for latent space alignment while mitigating catastrophic forgetting [26]. The second stage enhances perceptual fidelity through adversarial finetuning of the Low-rank adapted (LoRA [18]) latent U-Net. Finally, the third stage extends the framework to bursts by generalizing the classical align-and-merge operation to latent space. A lightweight spatio-temporal transformer [74] refines the center-frame latent using context from surrounding frames, thereby mitigating temporal artifacts such as flicker and content drift.

A key design goal of our method is to reconcile two complementary paradigms in image reconstruction. Classical quanta burst pipelines follow an align-and-merge strategy (e.g., QBP, QUIVER). In contrast, recent generative image restoration methods leverage large pretrained text-to-image models as strong natural-image priors but typically operate on a single frame. Our approach unifies these perspectives. Main contributions are as follows:

- A modular approach that adapts large-scale generative priors (e.g., Stable Diffusion [47]) to the extreme regime of quanta burst reconstruction.

- A learning-based method to jointly denoise, demosaic, and align bursts from color single-photon sensors, and a latent-space spatio-temporal transformer that enhances temporal consistency and mitigates content drift.
- The first real-world color SPAD burst dataset and a new *eXtreme motion + Deforming (XD)* video dataset.

This work takes a first step toward adapting internet-scale generative priors to photon-limited sensing with quanta cameras, enabling high-quality color and monochrome reconstructions under ultra high-speed motion.

## 2. Related Work

**Denoising for Conventional Cameras.** State-of-the-art denoising networks [7, 73] set strong baselines for conventional image restoration. NAFNet [7] adopts a minimalist residual design without nonlinear activations, prioritizing spatial fidelity and efficiency, while Restormer [73] employs transformer-based long-range modeling to capture complex noise characteristics. Although effective on Poisson-Gaussian noise typical of standard sensors, these models are not directly applicable to SPAD imagery, where photon shot noise, binary quantization, and extreme sparsity fundamentally alter the noise distribution. To provide representative baselines, we adapt and finetune both architectures for SPAD denoising and demosaicing.

**Quanta Burst Reconstruction.** Quanta Burst Photography

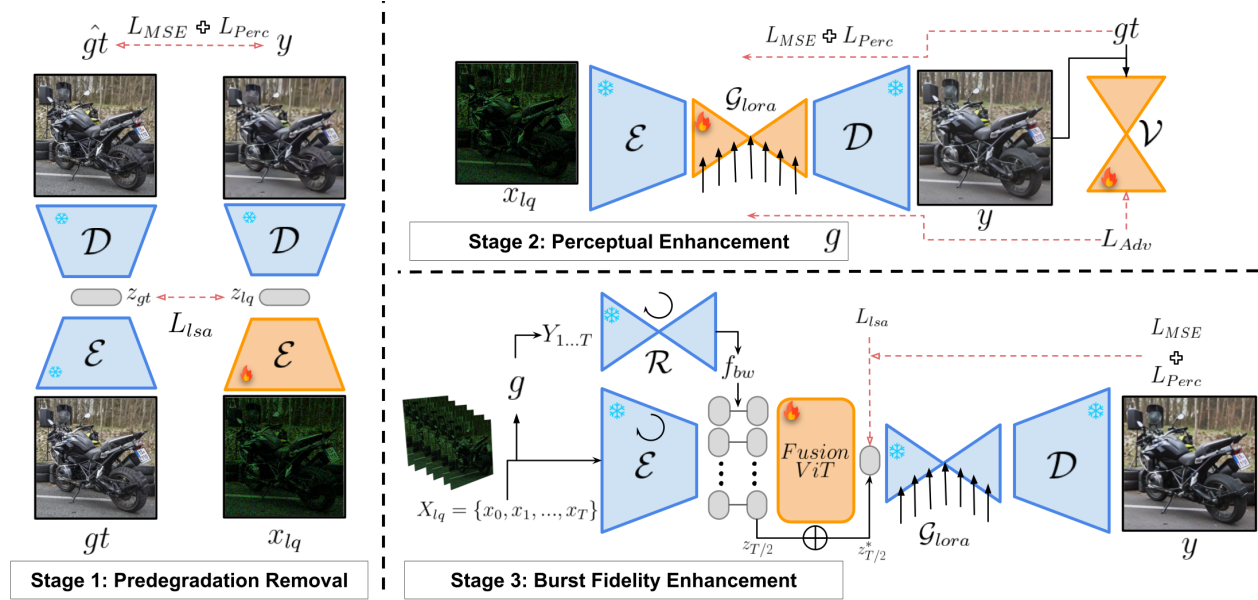


Figure 2. **Overview of gQIR.** Three-stage framework for quanta burst reconstruction: (S1) a quanta-aligned VAE for joint denoising and demosaicing of SPAD nano-bursts, (S2) an adversarially finetuned LoRA [18] latent U-Net initialized with Stable Diffusion [47] weights for perceptual enhancement, and (S3) a latent burst FusionViT for motion-aware spatio-temporal fusion of burst of nano-burst inputs.

(QBP) [78] introduced an offline burst-denoising pipeline combining block-matching temporal alignment with frame merging and Wiener filtering. A subsequent work [37] extended this framework to color, analyzing filter design for SPAD sensors and proposing a color-QBP variant to reconstruct RGB images from mosaiced bursts. Later reconstruction methods incorporated learned components while retaining the align-and-merge philosophy. QUIVER [8] uses an 11-frame window of 3-bit bursts, applying light pre-denoising to stabilize optical-flow estimation (SpyNet [46]) before recurrent fusion, while QuDI [9] replaces QUIVER’s recurrent denoiser with a time-conditioned U-Net, unrolling into a DDPM-like [17] formulation. Work on multi-bit quanta and QIS reconstruction [10, 11] and on binary-to-multi-bit mappings [35] also use learned models for photon-efficient sensing. In parallel, efficiency-oriented approaches [55, 76] reduce bandwidth by compressing quanta sequences before reconstruction, offering complementary benefits but addressing an orthogonal goal to fidelity. Taken together, existing methods remain task-specific and operate without largely-pretrained priors, a gap our approach addresses.

**Generative Image Restoration.** Diffusion-based generative models [17, 54] introduced iterative denoising to learn natural image distributions, and large text-to-image (T2I) diffusion models [6, 14, 32, 44, 47] have since become strong priors for a range of restoration tasks [13, 59, 60, 62, 63, 66, 72]. Recent blind restoration methods [19, 33, 34, 70] demonstrate that such priors can be adapted to diverse degradations while achieving high perceptual quality. These approaches typically align the latent

space of a pretrained diffusion model to a target degradation via lightweight finetuning or adapters. Unlike conventional diffusion sampling pipelines, our method performs deterministic single-step restoration. The generative prior is incorporated through LoRA-adapted diffusion weights and adversarial training, allowing reconstructions to follow the learned natural-image manifold without iterative sampling.

### 3. Methodology

Our architecture is structured as a 3-stage pipeline to progressively bridge the gap between sparse photon measurements and natural-image priors. Stage 1 produces measurement-consistent latent codes by adapting the VAE encoder to the statistics of Bernoulli photon detections. Stage 2 injects semantic knowledge from a pretrained diffusion model through a LoRA adaptation, improving perceptual quality. Stage 3 extends the reconstruction to burst by performing motion-aware alignment. This stage generalizes the classical align-and-merge paradigm of quanta burst reconstruction to a generative latent representation. This framework, described subsequently, is shown in Fig. 2.

#### 3.1. Image Formation Model

Photon imaging fundamentally differs from conventional cameras: each pixel registers discrete photon arrivals rather than analog intensities, with negligible read noise in SPADs. A physically consistent probabilistic rendering pipeline is used to synthesize a quanta observation from a clean sRGB ground truth image  $x_{gt} \in [0, 1]^{H \times W \times 3}$ . First,  $x_{gt}$  is mapped to linear radiance space via  $x_{lin} = x_{gt}^\gamma$ , using a fixed  $\gamma = 2.2$ , so that pixel intensities scale proportionally with scene irradiance. Given this linear image, a SPAD




Figure 3. **Qualitative comparison – single 3-bit frame reconstructions.** Conventional finetuned baselines over-smooth high-frequency structures, especially in distant depth planes and textured regions, whereas gQIR preserves sharper details and more faithful facial features, benefitting from the inclusion of FFHQ faces [24] in the training set.

records whether at least one photon arrives during the exposure. Assuming Poisson arrivals with rate  $\lambda$ , the SPAD output  $x_{spad}$  follows a Bernoulli distribution [78]:

$$x_{spad} = \text{Bern}(1 - e^{-\lambda}) = \text{Bern}(1 - e^{-(\alpha \cdot x_{lin})}), \quad (1)$$

where  $\lambda = \alpha \cdot x_{lin}$ , with  $\alpha$  controlling the expected photon-per-pixel (PPP) level. The expected PPP is  $\mathbb{E}[\lambda] = \alpha \mathbb{E}[x_{lin}]$ . In our setup,  $\alpha = 1.0$  or an average PPP of 3.5 matches the illumination levels in [8, 9].

To simulate a color SPAD, we apply a randomly sampled Bayer pattern  $\pi \in \{\text{RGGB}, \text{GRBG}, \text{BGGR}, \text{GBRG}\}$ , yielding a mosaiced binary frame  $x_{lq} = M_{\pi}(x_{spad})$ . An  $N$ -frame or  $\log_2(N + 1)$ -bit mosaiced observation is:

$$x_{lq} = \frac{1}{N} \sum_{i=1}^N M_{\pi}[\text{Bern}(1 - e^{-\alpha \cdot x_{lin}})] \quad (2)$$

### 3.2. Stage 1: Quanta Aligned VAE

Generative restoration methods [33, 34, 70] make their VAE’s encoders degradation-aware by optimizing the ob-

jective:  $\mathcal{L}_{\mathcal{E}_{\phi^*}} = \|\mathcal{D}(\mathcal{E}_{\phi^*}(x_{LQ})) - \mathcal{D}(\mathcal{E}_{\phi^*}(x_{GT}))\|_2^2$ , where  $\mathcal{D}$  denotes the frozen decoder,  $\mathcal{E}_{\phi^*}$  denotes the finetuned encoder, and  $x_{LQ}$ ,  $x_{GT}$  denote the low-quality and ground truth images respectively. This step is commonly known as degradation pre-removal [33, 70] which partially addresses restoration. However, naively applying this step leads to catastrophic latent-space forgetting in our case due to the extreme photon-shot noise in SPADs. The encoder eventually finds a smoothed shortcut solution to generate the same image, regardless of the input, as shown in Fig. 4. To address this, we introduce two key modifications: *deterministic mean encoding* and *latent space alignment loss*.

**Deterministic Mean Encoding.** Instead of stochastic sampling from the posterior  $q_{\phi}(z|x_{lq}) = \mathcal{N}(\mu_{\phi}(x_{lq}), \sigma_{\phi}^2(x_{lq}))$  we use the deterministic mean:  $\mathbb{E}_{q_{\phi}(z|x_{lq})} = \mu_{\phi}(x_{lq})$  obtained from the frozen, pre-trained encoder  $\mathcal{E}_{\phi}$ . This deterministic formulation avoids stochastic variance amplification, which is particularly important since  $x_{lq}$  is severely corrupted by photon-shot noise and already exhibits heavy-tailed statistics. Our objective is to maximize fidelity by

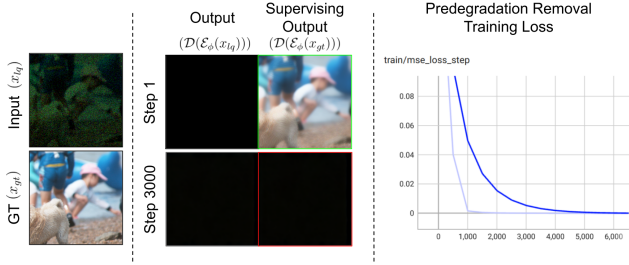


Figure 4. **Encoder collapse under predegradation removal loss** [34, 70]. The encoder  $\mathcal{E}_{\phi^*}$  learns a perceptually meaningless shortcut thus producing constant outputs. Since the trainable encoder controls both, the supervision and prediction terms, the training curve quickly converges to a degenerate optimum without our proposed modifications.

preserving the latent structure of the underlying clean scene. We achieve this by adding our new latent loss to modified reconstruction losses described as follows.

**Latent Space Alignment (LSA) Loss.** Since the decoder remains fixed throughout training to preserve its internet-scale training learned manifold, the encoder ( $\mathcal{E}_{\phi^*}$ ) must perform simultaneous denoising and demosaicing to produce a clean latent representation from  $x_{lq}$ . We enforce latent consistency between the low-quality and ground-truth embeddings using the following alignment loss:

$$\mathcal{L}_{lsa} = \|\mu_{\phi^*}(x_{lq}) - \mu_{\phi}(x_{gt})\|_2^2. \quad (3)$$

It is worth noting that the second term, unlike [34, 70] utilizes a frozen copy of the pre-trained Encoder  $E_{\phi}$ . This safeguards against the predegradation removal encoder’s collapse as shown in Fig. 4.

**Pixel Space Losses.** We also use MSE and LPIPS loss [75]:

$$\mathcal{L}_{MSE} = \|\mathcal{D}(\mu_{\phi^*}(x_{lq})) - \mathcal{D}(\mu_{\phi}(x_{gt}))\|_2^2, \quad (4)$$

$$\mathcal{L}_{perc} = \|\Phi(\mathcal{D}(\mu_{\phi^*}(x_{lq}))) - \Phi(\mathcal{D}(\mu_{\phi}(x_{gt})))\|_2^2, \quad (5)$$

where  $\Phi$  denotes a VGG-19 [52] backbone.

**Overall loss** is given by:

$$\mathcal{L}_{\mathcal{E}_{vae}} = \lambda_{lsa}\mathcal{L}_{lsa} + \lambda_{MSE}\mathcal{L}_{MSE} + \lambda_{perc}\mathcal{L}_{perc}, \quad (6)$$

where  $\lambda_{lsa}$ ,  $\lambda_{MSE}$  and  $\lambda_{perc}$  are scalar hyperparameters.

### 3.3. Stage 2: Perceptual Enhancement

Stage 1’s VAE alignment enables joint denoising and demosaicing, recovering structural, chromatic, and low-frequency details. In Stage 2, we finetune the pretrained diffusion backbone to refine the reconstruction, enhancing high-frequency details and improving perceptual quality.

Due to the extremely high data capture rate of SPAD sensors, reconstruction algorithms are often faced with huge amount of data processing, motivating the design of single-step algorithms. Adversarial training has been established as an effective way of distilling the diffusion prior to a single-step model [23, 32, 49, 68, 69]. Specifically, [34]




Figure 5. **Dynamic spatio-temporal latent burst merging.** Naive averaging of flow-aligned burst latents yields blur under scene motion. FusionViT instead adaptively weights latents by motion and proximity to the reference, producing a sharper output.

demonstrated a theoretical guarantee for stable GAN training by initializing the LoRA-initialized denoising network  $\mathcal{G}_{lora}$  with the prior’s diffusion weights  $\mathcal{U}_{\phi}$ , which ensures small initial gradients for a stable start of GAN training. We design a multilevel ConvNext-Large [36] backbone discriminator  $\mathcal{V}_{\theta}$  modified from [34] to adversarially train  $\mathcal{G}_{lora}$  using the standard min-max GAN objective [16]:

$$\min_{\phi} \max_{\theta} \mathbb{E}_{x \sim p_{x_{gt}}} [\log \mathcal{V}_{\theta}(x)] + \mathbb{E}_{x \sim p_{x_{lq}}} [\log(1 - \mathcal{V}_{\theta}(\mathcal{G}(x)))]. \quad (7)$$

The generator is additionally updated [29] by the pixel space reconstruction and the perceptual loss. Overall:

$$\mathcal{L}_{\mathcal{G}_{lora}} = \mathcal{L}_{adv} + L_{perc} + \|\mathcal{D}(\mathcal{G}_{lora}(\mu_{\phi^*}(x_{lq}))) - x_{gt}\|_2^2. \quad (8)$$

### 3.4. Stage 3: Latent Burst Imaging

Next, we extend our model with a burst window [8, 9, 78] to exploit temporal information in a quanta burst sequence. We generalize the classical align-and-merge philosophy of QBP [78] to the VAE’s latent space. We first compute optical flow to align all burst latent maps to the center latent map  $z_c$ . However, a pre-trained optical flow estimator  $\mathcal{R}$  cannot accurately warp the latent map despite the presence of rich semantic information. Moreover, applying pre-trained models directly to the low-quality sequence  $X = (x_{lq}^0, \dots, x_{lq}^i, x_{lq}^{i+1}, \dots)$  fails due to a significant domain gap (See supplementary). To bridge the gap, we first reconstruct all frames:  $Y = \mathcal{D}_{\phi}(\mathcal{G}_{lora}(\mathcal{E}_{\phi^*}(X_{lq})))$  and then use RAFT [57], pretrained on FlyingThings3D [39], to estimate the flow  $\mathcal{O} = \mathcal{R}(Y)$ , similar in spirit to pre-denoising [8, 9] or temporal aggregation [78].

Once optical flow is estimated, all burst frames are warped to the reference and merged. Naively averaging these aligned latents produces significant blur due to motion (Fig. 5). To overcome this, we introduce a pseudo-3D miniViT [74] ( $\mathcal{F}$ ) that applies sub-quadratic windowed attention across time and the spatial axis, enabling dynamic spatio-temporal burst fusion into a single high-fidelity latent code. Furthermore, the output of the FusionViT is modulated and residually added to the center latent  $z_{T/2}$  as shown in Fig. 2. The modulation is a learned scalar  $\delta$ , initialized at 0.05, to adaptively add or subtract the fused details prior to feeding the latent code into the generative network  $\mathcal{G}_{lora}$ .




Figure 6. **Qualitative comparison – burst reconstruction.** We simulate 1:1 GT–SPAD bursts by averaging 77 binary frames per input, preserving the original scene frame rate. QBP yields blurred reconstructions under fast motion due to small burst input while QUIVER breaks down due to motion-blurred nano-bursts, from realistic sampling. Our burst pipeline consistently recovers sharper structure and higher fidelity across extreme motion regimes, from 1000 to 100k fps.

We freeze all other networks and supervise FusionViT with the following overall loss, similar to Stage 1:

$$\mathcal{L}_{fusion} = \|\mathcal{F}(\mu_{\phi^*}(X_{lq})) - \mu_{\phi}(x_{gt})\|_2^2 + \|\mathcal{D}(\mathcal{G}_{lora}(\mathcal{F}(\mu_{\phi^*}(X_{lq})))) - x_{gt}\|_2^2 + \mathcal{L}_{perc} \quad (9)$$

### 3.5. Implementation Details

**GT-SPAD Sim.** We use  $\alpha=1.0$  for all simulations ( $\mathbb{E}[\text{PPP}]$ : 3.5) with random Bayer patterns. S1-2 use nano-bursts formed by averaging 7 i.i.d. binary realizations of 1 GT while Stage 3 uses 1 binary frame realized from 1 GT frame.

**Datasets** We train on 2.81 M images and 44,575 videos combined from diverse image [1, 24, 31, 50, 53] and video datasets [22, 41, 51, 77]. Testing also adds UDM [67], SPMC [56], our XD dataset to respective test-splits.

**Hyperparameters and Training.** Stage 1 is trained for 600k steps on  $8 \times \text{A100}$  with LR  $10^{-5}$ , batch size 8, and

scalars:  $\lambda_{lsa}=0.1$ ,  $\lambda_{MSE}=10^3$ ,  $\lambda_{perc}=2$ . Stage 2 runs for 100k iterations on a single RTX 4090 at  $256 \times 256$  with losses  $\lambda_{adv}=0.5$ ,  $\lambda_{MSE}=500$ ,  $\lambda_{perc}=5$ . S3 FusionViT trains for 20k steps using RAFT [57] pretrained on FlyingThings3D [39], with  $\lambda_{lsa}=1$ ,  $\lambda_{MSE}=1000$ ,  $\lambda_{perc}=7.5$ . All stages use Adam [27] optimizer with LR  $\eta = 10^{-5}$  and  $\beta=(0.9, 0.999)$ . All implementations are in PyTorch [43].

## 4. Experiments

**Baselines.** We establish single-frame denoising baselines by finetuning two representative RGB methods, NAFNet [7] and Restormer [73] since no other baselines exist for single quanta image reconstruction task, to the best of our knowledge. We do not finetune InstantIR [19] as it relies on a pre-trained generative prior and is designed for test-time unknown degradation removal. For the burst stage, we compare with quanta baselines: QBP [78] and

Table 1. **Fidelity and perceptual quality of 3-bit nano-burst input single RGB frame reconstruction.** Fine-tuned Restormer and NAFNet attain higher PSNR due to optimization for lower distortion [2], leading to oversmoothing, while gQIR achieves higher perceptual quality, consistent with visual results in Fig. 3.

Method	Monochrome						Color					
	Full-Reference			Non-Reference			Full-Reference			Non-Reference		
	PSNR↑	SSIM↑	LPIPS↓	ManIQA↑	ClipIQA↑	MUSIQ↑	PSNR↑	SSIM↑	LPIPS↓	ManIQA↑	ClipIQA↑	MUSIQ↑
InstantIR [19]	10.787	0.178	0.651	0.187	0.346	36.651	7.928	0.101	0.736	0.197	0.358	32.211
ft-Restormer [73]	<b>28.728</b>	0.816	0.294	0.262	0.435	40.439	26.433	0.739	0.388	0.235	0.395	36.026
ft-NAFNet [7]	28.276	0.830	<b>0.261</b>	0.300	0.473	39.129	<b>26.881</b>	0.757	<b>0.338</b>	0.251	0.431	36.732
(Ours) qVAE	28.180	<b>0.863</b>	0.327	0.299	0.487	44.895	26.280	<b>0.791</b>	0.435	0.272	0.432	38.613
(Ours) gQIR	27.281	0.839	0.318	<b>0.331</b>	<b>0.547</b>	<b>45.614</b>	25.480	0.766	0.361	<b>0.313</b>	<b>0.490</b>	<b>42.038</b>

QUIVER [8], with a burst size of 3-bit 11 frames.

**Metrics.** We evaluate all methods using full-reference: PSNR, SSIM [61], LPIPS [75] and non-reference metrics: ManIQA [65], ClipIQA [58] and MUSIQ [25]. For video (sliding burst windows) comparisons, we evaluate temporal consistency using flow-warping metric:  $E_{warp}$  [28].

**Test Datasets.** We use 2 different curated testsets. The single image reconstruction set consists of 334 images while the burst set consists of 11 100-frame videos from XVFI-test, I2-2000fps and XD-Dataset. We also evaluate our burst method on the entire test split of I2-2000fps.

#### 4.1. Quantitative Evaluation

**Single Image Comparisons.** Fidelity and perceptual scores, evaluated at the finetuned baselines’ 384<sup>2</sup> training resolution, are provided in Tab. 1. It is well established that a trade-off exists between perceptual quality and distortion [2]. While existing denoising methods optimize for lower distortion as evidenced by full-reference metrics, the input 3-bit frame contains extremely sparse information, and optimizing for distortion leads to oversmoothing. By leveraging a strong generative prior, gQIR achieves superior perceptual quality, reflected in no-reference metrics and consistent with the qualitative comparisons in Sec. 4.2.

**Burst Comparisons.** We curate a subset of our entire test set, grouping sequences by frame rate to evaluate methods across a wide spectrum, unlike prior works with fixed fps [8, 9]. Specifically, we select 11 sequences from XVFI [51] (1k fps), I2-2000fps [8] (2k fps), and XD (2k–100k fps). Tab. 2 shows that Burst-gQIR consistently outperforms the baselines, particularly on the challenging XD dataset, where QBP and QUIVER exhibit substantial performance drop.

We further evaluate our method on the full I2-2000fps [8] test set as shown in Tab. 3. Despite a minor domain gap between a PPP of 3.5 (ours) and 3.25 [8], our method surpasses the previous best (QuDI [9]) by +2.17 dB.

#### 4.2. Qualitative Evaluation

**Single Image Comparisons.** Monochrome and color SPAD reconstruction comparisons are shown in Fig. 3. InstantIR fails to produce high-quality reconstructions due

to the mismatch between Poisson–Gaussian and Bernoulli noise statistics. Fine-tuned Restormer and NAFNet yield reasonable but over-smoothed results, whereas the proposed gQIR restores fine details and enhances perceptual quality.

**Burst Comparisons.** Fig. 6 presents qualitative comparisons for burst reconstruction. While QUIVER [8] is trained and evaluated on motion-blur-free nano-bursts generated by inflating 11 GT frames into 77 binary sequences (sampling 7 binary frames per GT), we adopt a more realistic setting by sampling a single binary frame per GT, introducing motion blur in each nano-burst. Due to this domain gap, QUIVER fails to produce high-quality reconstructions. QBP, designed for bursts with hundreds of frames, yields blurry outputs under fast motion. In contrast, Burst-gQIR delivers sharp, high-fidelity reconstructions, effectively handling motion while preserving perceptual quality.



Figure 7. **SD3.5’s VAE solves the text problem [4, 5].** SD3.5 [14] uses a 4× larger latent space than SD2.1. This yields sharper high-frequency details and legible text drawing capabilities.

#### 4.3. Ablation Studies

**Stage 1 Design Choices.** We ablate the introduced combination of Latent Space Loss (LSA) and deterministic sampling for quanta frames and observe that LSA provides a critical gradient to the encoder for convergence in Tab. 4.

**VAE Scaling Solves the Text Problem [4, 5].** Increasing the VAE’s latent dimensionality, markedly improves capac-

Table 2. **Burst reconstruction fidelity under extreme motion.** Our method achieves superior scores due to cleaner flow processing and dynamic burst merging while keeping the traditional align-and-merge philosophy aided with a generative prior.

Test-set (GT fps)	QBP [78]			QUIVER [8]			Burst-gQIR		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
XVFI [51] (1000)	12.014	0.370	0.647	23.001	0.751	0.575	<b>25.822</b>	<b>0.712</b>	<b>0.434</b>
I2-2000fps [8] (2000)	16.043	0.549	0.468	25.060	0.874	0.366	<b>31.214</b>	<b>0.878</b>	<b>0.296</b>
XD (2k - 100k)	12.780	0.409	0.458	20.096	0.790	0.421	<b>30.331</b>	<b>0.895</b>	<b>0.316</b>
Cumulative:	13.380	0.448	0.496	22.429	0.814	0.429	<b>29.832</b>	<b>0.856</b>	<b>0.330</b>

Table 3. **Burst Fidelity on I2-2k benchmark.** Despite the PPP mismatch, our method reaches superior fidelity on I2-2k [8].

Method	PSNR↑	SSIM↑
EMVD [38]	20.019	0.587
FloRNN [30]	21.034	0.679
QBP [78]	21.548	0.703
Transform Denoise [3]	21.317	0.718
QUIVER [8]	26.214	0.790
QuDi [9]	28.641	0.811
(Ours) Burst-gQIR (3.25 PPP)	<b>30.811</b>	<b>0.868</b>

Table 4. **Ablation - Stage 1 Design Choices and Losses.** Our latent space alignment loss and deterministic sampling gives the highest fidelity in 1 epoch for joint denoising and demosaicing. Both components are critical for meaningful convergence and avoiding catastrophic forgetting shown in Fig. 4. (512<sup>2</sup>px)

Variants	PSNR↑	SSIM↑	ManIQA↑
w/o det. encoding (A)	20.56	0.435	0.167
w/o LSA loss (B)	10.39	0.222	0.139
w/o (A) and (B)	10.30	0.218	0.136
<b>Ours</b>	<b>24.78</b>	<b>0.665</b>	<b>0.194</b>

ity thereby enabling text synthesis. We show a 4× larger aligned SD3.5 [14] qVAE in Fig. 7 and the supplementary. **Fidelity, Perceptualness and Video Stability.** We compare all three stages of our method based on fidelity, perceptual quality, and video stability over the video test set used in Fig. 6. Stage 2 enhances photorealism but incurs a higher degree of content drift. This is attributed to its greater emphasis on perceptualness during training. In contrast, Stage 3 is optimized for combining temporal information for higher fidelity. This naturally mitigates content drift, as demonstrated in Tab. 5. Qualitative video reconstruction comparisons are provided in the supplementary.

#### 4.4. Real World Testing.

gQIR reconstructs photorealistic images from real color SPAD captures without explicit correction for dark count or hot pixel as shown in Fig. 8; the only post-processing applied is gray-world white balancing. Interestingly, gQIR retains fidelity to the vignetting artifact inherent to our sensor prototype. More qualitative results on our real-world

Table 5. **Ablation: All stages – fidelity versus temporal stability.** Stage 2 improves fidelity over Stage 1 but slightly increases content drift, while Stage 3 provides the best overall trade-off between reconstruction quality and temporal stability. (256<sup>2</sup>px)

Stage	PSNR↑	SSIM↑	E*↓
Alignment (S1)	20.038	0.759	9.088
Perceptual (S2)	24.114	0.846	8.508
Fidelity (S3)	<b>27.630</b>	<b>0.869</b>	<b>8.005</b>




Figure 8. **Real color SPAD reconstructions.** Qualitative results on binary bursts captured with a 1Mpx passive color SPAD prototype at 6k fps. Insets show demosaicing via sum-and-average.

acquisitions are provided in the supplementary material.

## 5. Limitations and Outlook

This work presents the first use of large-scale generative priors for quanta burst reconstruction, introducing techniques tailored to emerging color SPAD sensors. Despite resolution scalability via VAE tiling, several limitations remain. Motion cues from Stage 2 can degrade under subtle inter-frame drift, suggesting that video-level or multi-frame diffusion priors may further improve temporal coherence. Second, our training assumes a fixed 3.5 PPP, which limits robustness under extremely low-light ( $PPP < 1$ ). Explicitly modeling PPP as a conditioning signal may enhance generalization across lighting and sensor characteristics. Third, the 8-bit limit of the pretrained VAE decoder restricts the native HDR of SPADs [20, 21]; developing HDR-capable decoders is an important next step.

## 6. Acknowledgments

This research was supported by the National Science Foundation (CAREER #1943149), the Office of Naval Research (N000142412155), and by a Sony Faculty Innovation Award. We also thank Ubicept for their valuable input while writing this paper.

## References

- [1] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1122–1131, 2017. 6
- [2] Yochai Blau and Tomer Michaeli. The perception-distortion tradeoff. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6228–6237, 2018. 7
- [3] Stanley H. Chan, Omar A. Elgendy, and Xiran Wang. Images from bits: Non-iterative image reconstruction for quanta image sensors. *Sensors*, 16(11), 2016. 8
- [4] Jingye Chen, Yupan Huang, Tengchao Lv, Lei Cui, Qifeng Chen, and Furu Wei. Textdiffuser: Diffusion models as text painters. *arXiv preprint arXiv:2305.10855*, 2023. 7
- [5] Jingye Chen, Yupan Huang, Tengchao Lv, Lei Cui, Qifeng Chen, and Furu Wei. Textdiffuser-2: Unleashing the power of language models for text rendering. In *Computer Vision – ECCV 2024: 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part V*, page 386–402, Berlin, Heidelberg, 2024. Springer-Verlag. 7
- [6] Junsong Chen, Jincheng YU, Chongjian GE, Lewei Yao, Enze Xie, Zhongdao Wang, James Kwok, Ping Luo, Huchuan Lu, and Zhenguo Li. Pixart- $\alpha$ : Fast training of diffusion transformer for photorealistic text-to-image synthesis. In *The Twelfth International Conference on Learning Representations*, 2024. 1, 3
- [7] Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. Simple baselines for image restoration. *arXiv preprint arXiv:2204.04676*, 2022. 2, 6, 7
- [8] Prateek Chennuri, Yiheng Chi, Enze Jiang, GM Dilshan Godaliyadda, Abhiram Gnanasambandam, Hamid R Sheikh, Istvan Gyongy, and Stanley H Chan. Quanta video restoration. In *European Conference on Computer Vision*, pages 152–171. Springer, 2024. 1, 3, 4, 5, 7, 8
- [9] Prateek Chennuri, Dongdong Fu, and Stanley H. Chan. Quanta diffusion, 2025. 1, 3, 4, 5, 7, 8
- [10] Yiheng Chi, Abhiram Gnanasambandam, Vladlen Koltun, and Stanley H. Chan. Dynamic low-light imaging with quanta image sensors. In *Computer Vision – ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI*, page 122–138, Berlin, Heidelberg, 2020. Springer-Verlag. 3
- [11] Joon Hee Choi, Omar A. Elgendy, and Stanley H. Chan. Image reconstruction for quanta image sensors using deep neural networks. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6543–6547, 2018. 3
- [12] Ming Ding, Wendi Zheng, Wenyi Hong, and Jie Tang. Cogview2: faster and better text-to-image generation via hierarchical transformers. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2022. Curran Associates Inc. 1
- [13] Linwei Dong, Qingnan Fan, Yihong Guo, Zhonghao Wang, Qi Zhang, Jinwei Chen, Yawei Luo, and Changqing Zou. Tsd-sr: One-step diffusion with target score distillation for real-world image super-resolution. *arXiv preprint arXiv:2411.18263*, 2024. 3
- [14] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English, and Robin Rombach. Scaling rectified flow transformers for high-resolution image synthesis. In *Proceedings of the 41st International Conference on Machine Learning*. JMLR.org, 2024. 1, 3, 7, 8
- [15] Eric R. Fossum. Modeling the performance of single-bit and multi-bit quanta image sensors. *IEEE Journal of the Electron Devices Society*, 1(9):166–174, 2013. 1
- [16] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2*, page 2672–2680, Cambridge, MA, USA, 2014. MIT Press. 5
- [17] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, pages 6840–6851. Curran Associates, Inc., 2020. 3
- [18] Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. 2, 3
- [19] Jen-Yuan Huang, Haofan Wang, Qixun Wang, Xu Bai, Hao Ai, Peng Xing, and Jen-Tse Huang. Instantir: Blind image restoration with instant generative reference. *arXiv preprint arXiv:2410.06551*, 2024. 1, 3, 6, 7
- [20] Atul Ingle, Andreas Velten, and Mohit Gupta. High flux passive imaging with single photon sensors. In *Proc. CVPR*, 2019. 8
- [21] Atul Ingle, Trevor Seets, M Buttafava, Shantanu Gupta, Alberto Tosi, Andreas Velten, and Mohit Gupta. Passive inter-photon imaging. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021. 8
- [22] Sacha Jungerman, Max Leblang, Shantanu Gupta, and Kaushtubh Sadekar. visionsim. <https://github.com/WISION-Lab/visionsim>, 2025. 6
- [23] Minguk Kang, Richard Zhang, Connelly Barnes, Sylvain Paris, Suha Kwak, Jaesik Park, Eli Shechtman, Jun-Yan Zhu, and Taesung Park. Distilling Diffusion Models into Conditional GANs. In *European Conference on Computer Vision (ECCV)*, 2024. 5
- [24] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4396–4405, 2018. 4, 6
- [25] Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. Musiq: Multi-scale image quality transformer.

- 2021 *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5128–5137, 2021. 7
- [26] Ronald Kemker, Marc McClure, Angelina Abitino, Tyler L. Hayes, and Christopher Kanan. Measuring catastrophic forgetting in neural networks. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence*. AAAI Press, 2018. 2
- [27] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. 6
- [28] Wei-Sheng Lai, Jia-Bin Huang, Oliver Wang, Eli Shechtman, Ersin Yumer, and Ming-Hsuan Yang. Learning blind video temporal consistency. In *European Conference on Computer Vision*, 2018. 7
- [29] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew P. Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. Photo-realistic single image super-resolution using a generative adversarial network. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 105–114, 2016. 5
- [30] Junyi Li, Xiaohe Wu, Zhenxing Niu, and Wangmeng Zuo. Unidirectional video denoising by mimicking backward recurrent modules with look-ahead forward ones. In *Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XVIII*, page 592–609, Berlin, Heidelberg, 2022. Springer-Verlag. 8
- [31] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2017. 6
- [32] Shanchuan Lin, Anran Wang, and Xiao Yang. Sdxl-lightning: Progressive adversarial diffusion distillation, 2024. 1, 3, 5
- [33] Xinqi Lin, Jingwen He, Ziyang Chen, Zhaoyang Lyu, Bo Dai, Fanghua Yu, Yu Qiao, Wanli Ouyang, and Chao Dong. Diffbir: Toward blind image restoration with generative diffusion prior. In *Computer Vision – ECCV 2024*, pages 430–448, Cham, 2025. Springer Nature Switzerland. 1, 3, 4
- [34] Xinqi Lin, Fanghua Yu, Jinfan Hu, Zhiyuan You, Wu Shi, Jimmy S. Ren, Jinjin Gu, and Chao Dong. Harnessing diffusion-yielded score priors for image restoration, 2025. 1, 3, 4, 5
- [35] Yehe Liu, Alexander Krull, Hector Basevi, Ales Leonardis, and Michael W. Jenkins. bit2bit: 1-bit quanta video reconstruction via self-supervised photon prediction. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. 3
- [36] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. *CoRR*, abs/2201.03545, 2022. 5
- [37] Sizhuo Ma, Varun Sundar, Paul Mos, Claudio Brushini, Edoardo Charbon, and Mohit Gupta. “seeing photons in color”. *ACM Transactions on Graphics (TOG)*, 2023. 1, 3
- [38] Matteo Maggioni, Yibin Huang, Cheng Li, Shuai Xiao, Zhongqian Fu, and Fenglong Song. Efficient multi-stage video denoising with recurrent spatio-temporal fusion. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3465–3474, 2021. 8
- [39] N. Mayer, E. Ilg, P. Häusser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. arXiv:1512.02134. 5, 6
- [40] K. Morimoto, J. Iwata, M. Shinohara, H. Sekine, A. Abdelghafar, H. Tsuchiya, Y. Kuroda, K. Tojima, W. Endo, Y. Maehashi, Y. Ota, T. Sasago, S. Maekawa, S. Hikosaka, T. Kanou, A. Kato, T. Tezuka, S. Yoshizaki, T. Ogawa, K. Uehira, A. Ehara, F. Inui, Y. Matsuno, K. Sakurai, and T. Ichikawa. 3.2 megapixel 3d-stacked charge focusing spad for low-light imaging and depth sensing. In *2021 IEEE International Electron Devices Meeting (IEDM)*, pages 20.2.1–20.2.4, 2021. 1
- [41] Seungjun Nah, Sungyong Baik, Seokil Hong, Gyeongsik Moon, Sanghyun Son, Radu Timofte, and Kyoung Mu Lee. Ntire 2019 challenge on video deblurring and super-resolution: Dataset and study. In *CVPR Workshops*, 2019. 6
- [42] C. Niclass, A. Rochas, P.-A. Besse, and E. Charbon. Design and characterization of a cmos 3-d image sensor based on single photon avalanche diodes. *IEEE Journal of Solid-State Circuits*, 40(9):1847–1854, 2005. 1
- [43] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. *PyTorch: an imperative style, high-performance deep learning library*. Curran Associates Inc., Red Hook, NY, USA, 2019. 6
- [44] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis, 2023. 1, 3
- [45] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *ArXiv*, abs/2204.06125, 2022. 1
- [46] Anurag Ranjan and Michael J. Black. Optical flow estimation using a spatial pyramid network. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2720–2729, 2016. 3
- [47] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2, 3
- [48] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Lit, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Raphael Gontijo-Lopes, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. Photorealistic text-to-image

- diffusion models with deep language understanding. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2022. Curran Associates Inc. 1
- [49] Axel Sauer, Dominik Lorenz, Andreas Blattmann, and Robin Rombach. Adversarial diffusion distillation. In *European Conference on Computer Vision (ECCV)*, pages 87–103, Cham, 2025. Springer Nature Switzerland. 5
- [50] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. Laion-5b: an open large-scale dataset for training next generation image-text models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2022. Curran Associates Inc. 6
- [51] Hyeonjun Sim, Jihyong Oh, and Munchurl Kim. Xvfi: extreme video frame interpolation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021. 6, 7, 8
- [52] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition, 2015. 5
- [53] Ivan Skorokhodov, Grigorii Sotnikov, and Mohamed Elhoseiny. Aligning latent and image spaces to connect the unconnectable. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14144–14153, 2021. 6
- [54] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv:2010.02502*, 2020. 3
- [55] Varun Sundar, Matt Dutson, Andrei Ardelean, Claudio Bruschi, Edoardo Charbon, and Mohit Gupta. Generalized event cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [56] Xin Tao, Hongyun Gao, Renjie Liao, Jue Wang, and Jiaya Jia. Detail-revealing deep video super-resolution. In *The IEEE International Conference on Computer Vision (ICCV)*, 2017. 6
- [57] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *Computer Vision – ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II*, page 402–419, Berlin, Heidelberg, 2020. Springer-Verlag. 5, 6
- [58] Jianyi Wang, Kelvin CK Chan, and Chen Change Loy. Exploring clip for assessing the look and feel of images. In *AAAI*, 2023. 7
- [59] Jianyi Wang, Zongsheng Yue, Shangchen Zhou, Kelvin C.K. Chan, and Chen Change Loy. Exploiting diffusion prior for real-world image super-resolution. 2024. 3
- [60] Yufei Wang, Wenhan Yang, Xinyuan Chen, Yaohui Wang, Lanqing Guo, Lap-Pui Chau, Ziwei Liu, Yu Qiao, Alex C Kot, and Bihan Wen. Sinsr: diffusion-based image super-resolution in a single step. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 25796–25805, 2024. 3
- [61] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004. 7
- [62] Rongyuan Wu, Lingchen Sun, Zhiyuan Ma, and Lei Zhang. One-step effective diffusion network for real-world image super-resolution. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. 1, 3
- [63] Rongyuan Wu, Tao Yang, Lingchen Sun, Zhengqiang Zhang, Shuai Li, and Lei Zhang. Seesr: Towards semantics-aware real-world image super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 25456–25467, 2024. 3
- [64] Rui Xie, Yinhong Liu, Penghao Zhou, Chen Zhao, Jun Zhou, Kai Zhang, Zhenyu Zhang, Jian Yang, Zhenheng Yang, and Ying Tai. Star: Spatial-temporal augmentation with text-to-video models for real-world video super-resolution, 2025. 1
- [65] Sidi Yang, Tianhe Wu, Shuwei Shi, Shanshan Lao, Yuan Gong, Mingdeng Cao, Jiahao Wang, and Yujiu Yang. Maniqa: Multi-dimension attention network for no-reference image quality assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1191–1200, 2022. 7
- [66] Tao Yang, Rongyuan Wu, Peiran Ren, Xuansong Xie, and Lei Zhang. Pixel-aware stable diffusion for realistic image super-resolution and personalized stylization. In *The European Conference on Computer Vision (ECCV) 2024*, 2023. 3
- [67] Peng Yi, Zhongyuan Wang, Kui Jiang, Zhenfeng Shao, and Jiayi Ma. Multi-temporal ultra dense memory network for video super-resolution. *IEEE Transactions on Circuits and Systems for Video Technology*, 2019. 6
- [68] Tianwei Yin, Michaël Gharbi, Taesung Park, Richard Zhang, Eli Shechtman, Fredo Durand, and William T Freeman. Improved distribution matching distillation for fast image synthesis. In *NeurIPS*, 2024. 5
- [69] Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Frédo Durand, William T Freeman, and Taesung Park. One-step diffusion with distribution matching distillation. In *CVPR*, 2024. 5
- [70] Fanghua Yu, Jinjin Gu, Zheyuan Li, Jinfan Hu, Xiangtao Kong, Xintao Wang, Jingwen He, Yu Qiao, and Chao Dong. Scaling up to excellence: Practicing model scaling for photo-realistic image restoration in the wild, 2024. 1, 3, 4, 5
- [71] Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, Ben Hutchinson, Wei Han, Zarana Parekh, Xin Li, Han Zhang, Jason Baldridge, and Yonghui Wu. Scaling autoregressive models for content-rich text-to-image generation. *Transactions on Machine Learning Research*, 2022. Featured Certification. 1
- [72] Zongsheng Yue, Jianyi Wang, and Chen Change Loy. Resshift: Efficient diffusion model for image super-resolution by residual shifting. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. 1, 3
- [73] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *CVPR*, 2022. 2, 6, 7
- [74] Jinnian Zhang, Houwen Peng, Kan Wu, Mengchen Liu, Bin Xiao, Jianlong Fu, and Lu Yuan. Minivit: Com-

- pressing vision transformers with weight multiplexing. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12135–12144, 2022. 2, 5
- [75] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 5, 7
- [76] Tianyi Zhang, Matthew Dutson, Vivek Boominathan, Mohit Gupta, and Ashok Veeraraghavan. Streaming quanta sensors for online, high-performance imaging and vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(3):1564–1577, 2025. 3
- [77] Shangchen Zhou, Peiqing Yang, Jianyi Wang, Yihang Luo, and Chen Change Loy. Upscale-A-Video: Temporal-consistent diffusion model for real-world video super-resolution. In *CVPR*, 2024. 6
- [78] Sizhuo “Ma, Shantanu Gupta, Arin C. Ulku, Claudio Brushini, Edoardo Charbon, and Mohit” Gupta. “quanta burst photography”. *ACM Transactions on Graphics (TOG)*, “39”(“4”), “2020”. 1, 3, 4, 5, 6, 8