

## Efficient neural supersampling on a novel gaming dataset

Antoine Mercier Ruan Erasmus Yashesh Savani Manik Dhingra Fatih Porikli Guillaume Berger  
Qualcomm AI Research\*

{amercier, rerasmus, ysavani, manidhin, fporikli, guilberg}@qti.qualcomm.com

### Abstract

*Real-time rendering for video games has become increasingly challenging due to the need for higher resolutions, framerates and photorealism. Supersampling has emerged as an effective solution to address this challenge. Our work introduces a novel neural algorithm for supersampling rendered content that is  $4\times$  more efficient than existing methods while maintaining the same level of accuracy. Additionally, we introduce a new dataset which provides auxiliary modalities such as motion vectors and depth generated using graphics rendering features like viewport jittering and mipmap biasing at different resolutions. We believe that this dataset fills a gap in the current dataset landscape and can serve as a valuable resource to help measure progress in the field and advance the state-of-the-art in super-resolution techniques for gaming content.*

### 1. Introduction

Real-time rendering has become increasingly difficult for video games due to the demand for higher resolutions, framerates and photorealism. One solution that has recently emerged to address this challenge consists in rendering at lower resolution and then use an upscaling technique to achieve the desired resolution. However, developing efficient upscaling solutions that balance speed and accuracy remains a challenge. Recently, several commercial solutions have been developed for gaming super-resolution, including those that are based on deep learning (DL) such as Nvidia’s DLSS [36] or Intel’s XeSS [11], as well as solutions that do not rely on machine learning, such as AMD’s FSR [18, 19]. Despite the availability of these commercial solutions, there has been relatively little published research on the application of DL-based super-resolution for gaming. We believe that one of the reasons why DL-based super-resolution for gaming has received little attention compared to super-resolution of natural content is that there is currently no standard, publicly available dataset for develop-

\*Qualcomm AI Research is an initiative of Qualcomm Technologies, Inc. and/or its subsidiaries.




Figure 1. Example images produced by our solution using neural networks of different sizes. These models produce 1080p outputs in respectively 1.08 ms, 1.53 ms, and 3.52 ms on an RTX 3090, which is  $4\times$  to  $12\times$  faster than previous work by Xiao *et al.* [59].



Figure 2. Example of data modalities available in the QRISP dataset. *First row, from left to right:* Native 270p, Negative 2 mipmap biased 270p, Negative 1.58 mipmap biased 360p, Negative 1 mipmap biased 540p. *Second row, from left to right:* 540p depth, 540p motion vectors, Native 1080p, Enhanced 1080p

ing gaming-specific super-resolution solutions. Researchers and developers who want to study or improve upon existing methods must create their own datasets, which can be a time-consuming and resource-intensive process.

Our work makes the following contributions:

- we release a dataset specifically designed for the research and development of gaming super-resolution algorithms. We show that models trained on this dataset




Figure 3. Example of color, depth and motion vector images from the QRISP dataset. In total, our dataset contains 8760 frames (7260 for training, 1500 for testing) from 13 distinct scenes, rendered at different resolutions ranging from 270p to 1080p. More details can be found in the supplementary materials.

can compete and outperform the quality levels obtained by commercial solutions such as DLSS [36].

- we propose an efficient gaming super-resolution architecture which leverages auxiliary modalities (subpixel accurate motion vectors, depth) and graphics rendering features (viewport jittering, mipmap biasing) commonly-used for temporal anti-aliasing. Our solution is  $4\times$  more efficient than previous published work [59] for the same level of accuracy.

Overall, we believe that this work provides a new resource to measure progress in the field and help advance the state-of-the-art in gaming super-resolution.

## 2. Related work

**Generic super-resolution.** In recent years, DL-based approaches for super-resolution of natural content have become increasingly popular [13, 14, 52, 25, 38, 34, 37, 48], yielding state-of-the-art visual quality compared to interpolation and other algorithmic solutions. In this work, we focus mainly on approaches that exploit information gathered from consecutive frames, as multi-frame super-resolution (also called temporal supersampling in the gaming field) has become the de facto standard for video gaming [36, 11, 19]. Specifically, we consider online super-resolution architectures that can be efficiently stepped forward, as offline video enhancement approaches based on bidirectional mechanisms [24, 53, 9, 32] or sliding windows of input frames [8, 57, 56, 31] are not suitable for gaming applications. Efficient online multi-frame super-resolution is often based on recurrent convolutional architectures, either with explicit motion compensation [49] or without [16, 27, 26]. Alternatives to explicit motion compensation include architectures based on deformable convolutions [57], transformers [51, 1, 33] or dynamic upsam-

pling filters [28]. In gaming, however, explicit motion compensation is usually preferred, as the game engine can provide precise motion vectors and the neural network can be made much smaller if it doesn't have to learn how to copy past information over long distances.

**Gaming supersampling.** Temporal Anti-Aliasing (TAA) [62, 61, 29] and its upscaling forms [22, 17, 19] exploit samples from past frames to recover missing details in the current frame. Compared to single-frame anti-aliasing techniques [3, 39, 46], TAA has gained popularity over the past decade as it provides a good trade-off between accuracy and practicality, even in the context of deferred rendering where MSAA [3] becomes bandwidth prohibitive. A typical TAA pipeline includes a re-projection step [44, 50] to re-align historical data using accurate motion vectors, a history validation step to reject or rectify past samples that are invalid or stale due to *e.g.* occlusion, lighting or shading changes, and a blending (or accumulation) step to produce the final output. While TAA-based super-resolution approaches such as FSR2 [19] leverage hand-engineered heuristics to perform history validation and accumulation, DLSS [36], XeSS [11] and Xiao *et al.*'s work [59] have showed that these steps can be replaced by a neural net. In the rest of the paper, we compare our algorithm mainly against Xiao *et al.*'s [59], as the implementation details of DLSS and XeSS are not publicly available and, therefore, not reproducible.

**Graphics features traditionally used with supersampling.** Viewport jittering and negative mipmap biasing are two rendering techniques that are traditionally used to boost super-resolution accuracy. Viewport jittering consists in applying a sub-pixel shift to the camera sampling grid, and it is most useful when the camera is stationary as it ensures that consecutive frames contain complementary information about the scene. The subpixel offset typically follows

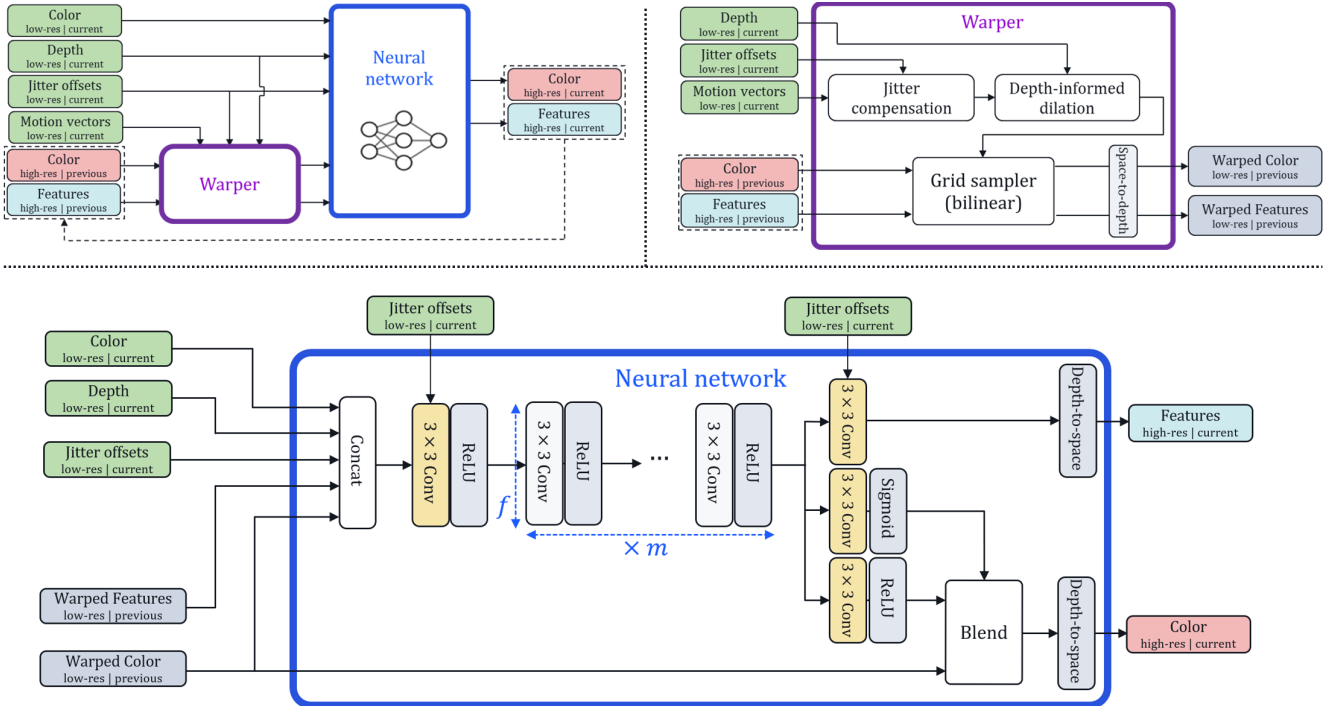


Figure 4. High-level overview of our multi-frame supersampling approach (*top-left*) and detailed description of its individual components: a warping module (*top-right*) and a reconstruction neural network (*bottom*).  $m$  and  $f$  refer to the number of intermediate conv layers and the number of features in these layers, respectively.

a fixed-length sequence parameterized by, for example, a Halton sequence. Negative mipmap biasing, on the other hand, reduces the amount of low-pass prefiltering applied to textures, resulting in low-resolution renders with more high-frequency details.

**Related datasets** While there are many datasets for single-frame [2, 6, 63, 40, 23] and video [43, 60] super-resolution of natural content, there is no publicly available dataset for gaming super-resolution. To the best of our knowledge, among existing datasets, Sintel [7] would be the closest candidate as it consists of synthetic images and provides motion vectors. It is however available at only one resolution, which is problematic because DL-based super-resolution models trained to reconstruct images from artificially downsized images tend to overfit the degradation method [35]. Besides, Sintel does not provide jittered or mipbiased samples, two key ingredients for gaming super-sampling. Our dataset is also significantly larger than Sintel ( $5\times$  more frames, and available at higher resolution).

### 3. The Qualcomm Rasterized Images for Super-resolution Processing dataset

The *Qualcomm Rasterized Images for Super-resolution Processing* dataset, abbreviated later in this paper as QRISP, was specifically designed to facilitate the development and research of super-resolution algorithms for gaming appli-

cations. To the best of our knowledge, this dataset has no publicly available equivalent.

**Data modalities.** The dataset consists of sequences of rasterized images captured at 60 frames per second with multiple modalities rendered at different resolutions ranging from 270p to 1080p. For each frame, the dataset includes color, depth, and motion vectors with different properties such as mipmapbiasing, jittering, or both. Jittered samples are achieved by shifting the camera using a sub-pixel offset drawn from a cyclic Halton(2, 3) sequence of length 16 and we occasionally include stationary segments in the camera path to make models trained on this dataset more robust to the “static” scenario. For low-resolution renders, we disable MSAA or any other frame-blurring anti-aliasing techniques and adjust the texture mip levels using a negative offset set to  $-\log_2(S)$  where  $S$  is the per-dimension scaling factor, as typically done in gaming super-sampling [29, 36, 18, 19, 59]. For high-res images, we target high-quality 1080p color images which were obtained by 2x-downsizing 2160p renders with MSAAx8 applied, as done in [59]. Figure 2 shows an example of such an “enhanced” target image (see the bottom-right crop), along with the corresponding low-resolution renders.

**Scene diversity and composition.** The dataset is diverse, with a variety of backgrounds and models to enable better generalization to new video games. There are 13 scenes in

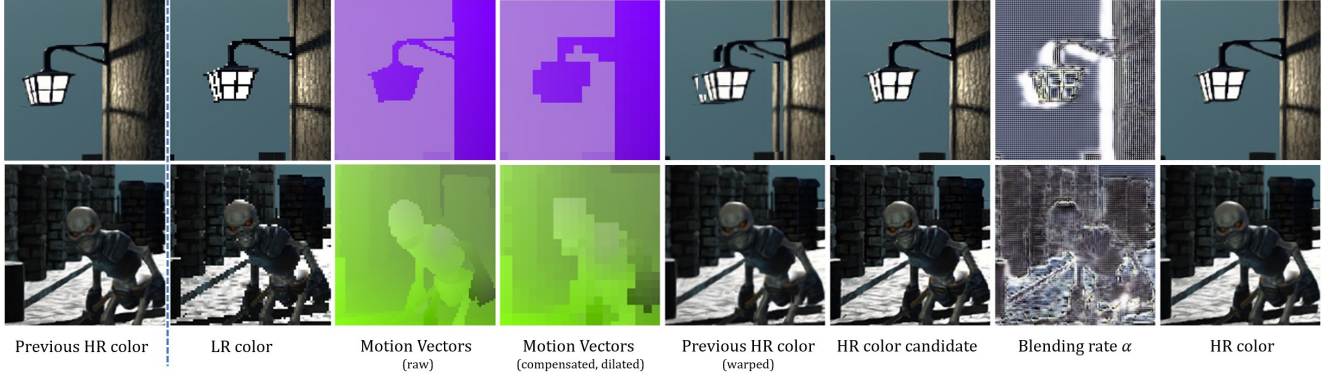


Figure 5. Visualization of data instances used at different steps of the algorithm (*from left to right*): the previous frame’s high-resolution output, the current low-resolution render, the motion vectors before and after pre-processing, the re-projected output from the previous frame, the new high-resolution color candidate for the current frame, the blending mask  $\alpha$ , and the final output for the current frame. Note that a low value of  $\alpha$  (dark) means that the color from the previous output is retained; a high value of  $\alpha$  (bright) means that the color from the previous timestep is discarded in favour of the new candidate color.




Figure 6. Visualization of the blending mask ( $8 \times 8$  crops) over consecutive timesteps along with the corresponding subpixel-shifted sampling locations (red dots) on a surface that does not need history rejection for  $2 \times$  (first row) and  $4 \times$  upscaling (second row). When the previous samples are still relevant for the current frame, the model tends to only update the pixels at the sampling location.

total, with 10 scenes allocated for training and the remaining 3 reserved for evaluation. Some of these scenes can be seen in Figure 3 and more samples can be found in the supplementary material. The data was generated using the Unity game engine [20], with 3D assets sourced either from the Unity Asset Store, or from open-source projects. The list of Unity assets used in this work can be found in the supplementary material. To make the data more representative of realistic gaming scenarios, animated characters were incorporated to the scene. We also added textual UI elements on top of animated characters to make the algorithms more robust to elements without associated depth or motion vector information.

**Commercial baselines.** In this dataset, we have also included images upscaled by commercial solutions integrated into Unity on the same frames used for evaluation. At the time of the dataset collection, these included Nvidia’s DLSS

2.2 and AMD’s FSR 1.2<sup>1</sup>, which can serve as reference baselines to assess the performance of new algorithms.

We believe that releasing this dataset will be beneficial for many, as it provides a time-saving alternative to extracting synchronized LR-HR pairs from a game engine, with additional modalities such as depth and motion vectors, and properties like jittering or mipmap biasing. This dataset was primarily created to advance the development of super-resolution algorithms for gaming applications, but we believe that it can also be useful for other tasks, such as optical flow estimation.

## 4. Proposed algorithm

Our proposed neural algorithm for gaming super-sampling consists of two components: a warping module for re-projecting historical data to the current frame, and a neural network for reconstructing the current image at target resolution. The reconstruction neural network blends a candidate image with the output image from the previous timestep using subpixel accurate motion vectors for motion compensation. Figures 4 and 5 provide an overview of the solution and a visualization of data instances at various steps of the algorithm, respectively.

### 4.1. Warping module

As seen in the top-right diagram of Figure 4, the motion compensation in our proposed algorithm is divided into three steps, which are described in detail below.

**Jitter compensation.** We remove the viewport jittering contribution to the motion vector values. This is achieved

<sup>1</sup>We do not compare against FSR 1.2 in this paper as we focus on multi-frame supersampling approaches.




Figure 7. Super-resolution results (2× and 3×) by our algorithm vs bicubic upscaling, Xia *et al.*'s approach [59] and DLSS 2.2 [36].

	Ours-S			Ours-M				Ours-L			Xiao <i>et al.</i>		
Scaling factor	2×	2×	3×	4×	2×	3×	4×	2×	3×	4×	2×	3×	4×
MV dilation	0.21	0.21	0.21	0.21	0.21	0.21	0.21	0.21	0.21	0.21	-	-	-
Warping	0.15	0.15	0.15	0.15	0.15	0.15	0.15	0.15	0.15	0.15	0.92	0.92	0.92
Neural network	0.72	1.17	0.81	0.70	3.16	1.74	1.26	13.01	12.92	12.88	13.01	12.92	12.88
<b>Total</b>	<b>1.08</b>	<b>1.53</b>	<b>1.17</b>	<b>1.06</b>	<b>3.52</b>	<b>2.10</b>	<b>1.62</b>	<b>13.93</b>	<b>13.84</b>	<b>13.80</b>			

Table 1. Profiling results (in milliseconds) of our 2×, 3× or 4× architectures vs Xiao *et al.* on an RTX 3090 at 1080p target resolution. These timings were all obtained using Nvidia’s TensorRT [12] for the neural network execution, using FP16 precision.

by adding the jitter offset at frame  $t - 1$  and subtracting the jitter offset at frame  $t$  from the motion vector at frame  $t$ :

$$MV_t = MV_t + J_{t-1} - J_t \quad (1)$$

**Depth-informed dilation.** This step modifies the motion vectors to reduce aliasing of foreground objects in re-projected images. This is achieved by producing a high-resolution block-based motion vector grid, where each block contains the motion vector value of the frontmost (i.e. lowest depth value) pixel within the block. In our experiments, we use a block size of 8×8 at high resolution. Similar ideas have been used in [29, 61, 19].

**Re-projection.** The preprocessed motion vectors are used to perform a bilinear warp and realign the previous

timestep’s high-resolution color images and neural features to the current frame. A space-to-depth operation is then applied to map the warping outputs to the input resolution, as in FRVSR [49].

## 4.2. Neural network

Our neural network architecture is similar to the efficient single-frame super-resolution architectures of [15, 5]. We use 3×3 Conv-ReLU blocks and a relatively small number of layers and channels. The output is mapped to high resolution using a depth-to-space operation. We however modify the architecture to make it suitable for multi-frame inputs and jittering:

**Additional inputs and outputs.** In addition to color information  $C_t$ , the neural network  $F$  takes as inputs depth  $D_t$ , the jitter offset  $J_t$ , the previous color output  $Y_{t-1}$  and

features  $f_{t-1}$  re-aligned by the warper  $W$ . It returns a pixel-wise blending mask  $\alpha$  obtained using a sigmoid activation, a high-resolution candidate color image  $\tilde{Y}_t$ , and recurrent features  $f_t$  for the next timestep:

$$\alpha, \tilde{Y}_t, f_t = F(C_t, D_t, J_t, W(Y_{t-1}), W(f_{t-1})) \quad (2)$$

**Blending.** The candidate color image returned by the neural net is combined with the previous output using  $\alpha$ :

$$Y_t = \alpha * \tilde{Y}_t + (1 - \alpha) * W(Y_{t-1}) \quad (3)$$

Figure 6 illustrates how the blending mask evolves over consecutive timesteps and shows that the model tends to use the candidate pixels located at the current sampling location but retains the previous samples everywhere else ( $\alpha \approx 0$ ). The neural net is also able to identify and discard samples ( $\alpha \approx 1$ ) from the re-projected color image that are outdated due to appearance changes or dis-occlusion (see Figure 5).

**Jitter-conditioned convolutions.** To facilitate alignment of the low-resolution color input, which is sub-pixel shifted due to jittering, we predict the kernel weights of the first and last convolution modules using an MLP conditioned on the jitter offset  $J_t$ . This is different from kernel prediction networks commonly used in denoising tasks [55, 64, 41] or burst image super-resolution [58, 10] where a separate kernel is predicted for each pixel, as we only predict one kernel for the entire frame, using a two-dimensional vector (i.e. the jitter offset for the current frame) as the only conditioning variable. At inference time, the jittering sequence is known in advance, so we pre-compute the kernel weights for each jitter offset and re-load the corresponding kernel whenever a new jittered frame is generated. This allows the neural network to more accurately realign the subpixel-shifted input data to the target resolution grid, with no additional computational overhead. The concurrent work of [21] used a similar technique for temporal anti-aliasing.

**Comparison to previous work.** Xiao *et al.*'s neural architecture consists of a feature extraction network that processes the current low-resolution frame, an upsampling module, a warping module that recursively realigns a rolling buffer of four upsampled feature maps, a feature reweighting network for history rejection, and a U-Net-style [47] reconstruction network to produce the final output. In comparison, our approach better leverages viewport jittering and has notable advantages in terms of speed and memory. First, all convolutional layers in our architecture run at the input resolution. Second, our historical data only consists of 4 high-resolution channels compared to 48 for Xiao *et al.*, resulting in a larger memory footprint and a higher latency for the re-alignment step.

Upscaling	Method	Metric		
		PSNR	SSIM	LPIPS
2×	Bicubic	29.51	0.8672	0.219
	DLSS 2	30.21	0.8816	0.187
	Xiao <i>et al.</i>	31.89	0.9075	0.140
	Ours-S (f16-m1)	31.18	0.8941	0.160
	Ours-M (f32-m3)	31.80	0.9044	0.140
	Ours-L (f64-m5)	<b>32.21</b>	<b>0.9115</b>	<b>0.134</b>
3×	Bicubic	27.61	0.8034	0.322
	Xiao <i>et al.</i>	30.24	0.8729	0.200
	Ours-M (f32-m3)	30.23	0.8655	0.203
	Ours-L (f64-m5)	<b>30.67</b>	<b>0.8747</b>	<b>0.187</b>
4×	Bicubic	26.42	0.7535	0.391
	Xiao <i>et al.</i>	29.02	0.8364	0.259
	Ours-M (f32-m3)	29.06	0.8305	0.258
	Ours-L (f64-m5)	<b>29.42</b>	<b>0.8403</b>	<b>0.238</b>

Table 2. PSNR, SSIM and LPIPS scores for our model, DLSS 2.2 [36] and our implementation of Xiao *et al.* [59] for 2×, 3× and 4× upscaling. For a more fine-grained analysis, a per-scene breakdown of PSNR and SSIM scores is available in the supplementary material.

## 5. Implementation details

We experiment with three variants of our models: *Ours-S (f16-m1)*, *Ours-M (f32-m3)* and *Ours-L (f64-m5)*, where *fX-mY* means that the architecture has  $Y$  intermediate conv layers and  $X$  feature channels. We adjust the number of recurrent features produced at low-resolution based on the scaling factor to end up with a single channel of features at high-resolution after depth-to-space. To predict the kernels of the first and last convs, we use a 7-layer MLP with 2048 hidden features and ReLU activations. Since our dataset uses the same fixed jittering sequence for the entire dataset, we also tried optimizing a set of 16 kernels, but it resulted in worse performance (0.1 dB PSNR drop on test scenes).

We use mini-batches of eight 16-frame clips with spatial resolution  $264 \times 264$  (at high resolution), an L1 loss, and train for 500k iterations using the Adam optimizer [30] with an initial learning rate of  $1e-4$ , decaying the learning rate by a factor 2 after 200k and 400k iterations. We optimize the models on 80% of the segments from each training scene and use the rest for validation.

## 6. Results analysis

### 6.1. Speed-accuracy tradeoff compared to existing solutions

Table 2 reports the average PSNR, SSIM and LPIPS scores obtained by our network, DLSS 2.2 and our implementation of Xiao *et al.*'s solution [59] on the test scenes.

Model variants	Entire test set		Static frames	
	PSNR	LPIPS	PSNR	Pixel Std
Baseline (f32-m3)	31.80	0.140	37.38	0.54
(-) MV dilation	31.97	0.136	37.47	0.54
(-) blending	31.80	0.144	36.86	1.10
(-) jitter	31.61	0.153	35.06	0.03
(-) MVs, (+) RAFT	31.09	0.169	35.77	1.21
(-) warping	30.69	0.183	37.35	0.53
(-) multiple frames	30.56	0.192	33.48	1.96
(-) first jitter conv	31.68	0.145	37.40	0.49
(-) last jitter conv	31.50	0.148	37.38	0.56
(-) jitter conv	31.26	0.156	37.20	0.56
(-2) training scenes	31.74	0.144	37.42	0.57
(-5) training scenes	31.78	0.142	37.14	0.76
(-8) training scenes	31.56	0.152	37.11	0.91
(+) perceptual loss	31.72	0.125	37.40	0.54

Table 3. Ablation study. We report the average PSNR and LPIPS scores on the entire test set, as well as PSNRs and average pixel-wise standard deviation on a static segment from the AbandonedSchool scene.




Figure 8. Visual impact of jittering on reconstructions of static scenes using our algorithm vs non-jittered alternatives.

Our small model outperforms DLSS, while our large architecture is slightly better than Xiao *et al.* These PSNR, SSIM and LPIPS improvements also manifest in visual quality improvements, as seen in Figure 7. We observe that DLSS 2.2 produces more ghosting artifacts (visible on the lantern in the first row or on the barrel in the third row) than the other approaches. The benefits from leveraging jittered samples are particularly visible in the reconstruction of static scenes (see Figure 8).

Table 1 demonstrates that our larger model generates 1080p images in 3.52 ms through  $2\times$  upscaling on an Nvidia RTX 3090, representing a  $4\times$  improvement compared to the architecture proposed by Xiao *et al.* [59] while maintaining the same level of accuracy. Our small architecture runs in 1.08 ms for the same workload. Additionally, our architecture scales better to larger magnification factors, with our  $4\times$  architecture offering an  $8.5\times$  speedup compared to Xiao *et al.* [59] for the same level of accuracy. We believe these timings could be improved using optimized CUDA kernels for the reprojection-related operations.

## 6.2. Ablation studies

In this section, we ablate individual components from our  $2\times$  medium-sized model, *Ours-M*, to illustrate the impact of each component on visual quality and stability.

**Depth-informed motion vector dilation.** While removing motion vector dilation improves PSNR and LPIPS scores (see Table 3), we found this step beneficial for the reconstruction of thin objects. This is visible in Figure 10 where we illustrate the effect of depth-informed dilation on the motion vectors, resulting in a warped image with less ghosting and aliasing artifacts, leading to a better reconstruction.

**Reconstruction quality and temporal stability on static scenes.** We evaluate the quality and temporal stability of model outputs on a section of the AbandonedSchool scene<sup>2</sup> where the camera is stationary. We report the average PSNR and pixel-wise standard deviation on these frames in Table 3, in addition to the average PSNR on the entire test set. We observe that:

- The single-frame variant of our architecture poorly reconstructs fine-grained details and is not temporally stable.
- Jittering is key to properly reconstruct static scenes: without it, the average PSNR on static scenes drops significantly. The benefits from leveraging jittered samples is also visible in Figure 8.
- Blending improves temporal stability. Without it, the pixel-wise standard deviation doubles (from 0.54 to 1.10) and we generally observe considerably more flickering artifacts with this variant.
- Temporal stability benefits from more (targeted) training data. The pixel-wise standard deviation score drops significant when the last 5 scenes are not used because only those contain static segments.

<sup>2</sup>Segment 0001, from frame 275 to 292

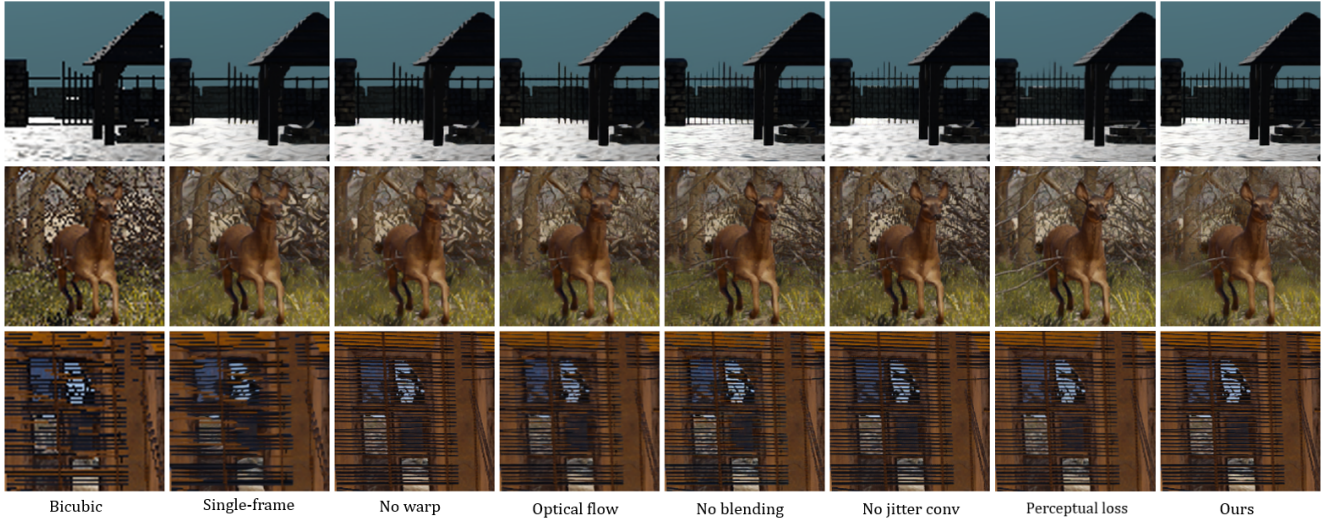


Figure 9. Ablation study using our  $2\times$  medium-sized architecture. We ablate the following components (*from left to right*): all components except the reconstruction network which we run on a single low-resolution frame, warping, motion vectors (which we replace by estimated optical flow obtained using RAFT [54]), jitter-conditioned convolutions and the L1 loss which we replace to the loss defined in [59].




Figure 10. Reconstruction of thin objects, with and without depth-informed motion vector dilation.

**On the benefits of jitter-conditioned convolutions** Table 3 shows the benefits of using jitter-conditioned kernels in the first and last convolution modules. Without these, we observe a 0.54 dB PSNR drop. Table 3 also suggests that the last layer is the one that benefits the most from using a jitter-conditioned kernel.

**On the importance of accurate motion compensation** To quantify the importance of accurate motion compensation, we trained our architecture both without motion compensation and on top of estimated motion vectors, which were estimated using RAFT [54] with weights pre-trained on Sintel [7]. In both cases, the results show a significant PSNR drop. When the camera remains static, the variant without motion compensation works well in terms of reconstruction and temporal stability (as seen in Table 3), but

the quality drops to the level of a single frame architecture when the camera moves (see the first two rows of Figure 9).

**Replacing the L1 loss to a perceptual loss** We find that the loss used in [59] (*i.e.* a weighted combination of SSIM and a VGG-based perceptual loss) improves the sharpness and perceived quality of high-frequency textures (most visible in the “deer” crop from Figure 9) at the cost of slightly more temporal inconsistencies. Quantitatively, this loss improves LPIPS scores significantly (from 0.140 to 0.125 for Ours-M) with minor PSNR and SSIM differences.

## 7. Conclusion and future work

In this work, we present a novel neural supersampling approach that is  $4\times$  more efficient than previous published work by Xiao *et al.* [59] while maintaining the same level of accuracy. We propose a new dataset specifically designed for the research and development of super-resolution algorithms for gaming applications. We believe this dataset fills a gap in the dataset landscape and can serve as a valuable resource to advance the state-of-the-art in super-resolution techniques for gaming content. When trained on the proposed dataset, our algorithm outperforms DLSS 2.2 [36] in terms of visual quality. In future work, we plan to investigate quantizing the neural network component of our approach as this has the potential to make the algorithm even more efficient [42, 4, 5, 45]. We also plan to explore whether more sophisticated losses can address the blurriness that sometimes arises in the reconstruction of severely under-sampled high-frequency textures (*e.g.* the grass and trees behind the deer in Figure 9).



## References

- [1] Davide Abati, Amir Ghodrati, Amirhossein Habibian, and Qualcomm AI Research. Efficient video super resolution by gated local self attention. In *2021 British Machine Vision Conference (BMVC)*, volume 2, page 3, 2021.
- [2] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 126–135, 2017.
- [3] Kurt Akeley. Reality engine graphics. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, pages 109–116, 1993.
- [4] Mustafa Ayazoglu. Extremely lightweight quantization robust real-time single-image super resolution for mobile devices. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2472–2479, 2021.
- [5] Guillaume Berger, Manik Dhingra, Antoine Mercier, Yashesh Savani, Ruan Erasmus, and Fatih Porikli. Quicksrnet: Plain single-image super-resolution architecture for faster inference on mobile platform. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition Workshops (MAI'23)*, 2023.
- [6] Marco Bevilacqua, Aline Roumy, Christine Guillemot, and Marie Line Alberi-Morel. Low-complexity single-image super-resolution based on nonnegative neighbor embedding. 2012.
- [7] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In A. Fitzgibbon et al. (Eds.), editor, *European Conf. on Computer Vision (ECCV)*, Part IV, LNCS 7577, pages 611–625. Springer-Verlag, Oct. 2012.
- [8] Jose Caballero, Christian Ledig, Andrew Aitken, Alejandro Acosta, Johannes Totz, Zehan Wang, and Wenzhe Shi. Real-time video super-resolution with spatio-temporal networks and motion compensation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4778–4787, 2017.
- [9] Kelvin CK Chan, Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Basicvsr: The search for essential components in video super-resolution and beyond. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4947–4956, 2021.
- [10] Wooyeong Cho, Sanghyeok Son, and Dae-Shik Kim. Weighted multi-kernel prediction network for burst image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 404–413, 2021.
- [11] Hisham Chowdhury, Rense Robert Kawiak, Gabriel Ferreira de Boer, and Lucas Xavier. Intel xess-an ai based super sampling solution for real-time rendering.(2022). In *Game Developers Conference*, 2022.
- [12] Nvidia Corporation. tensorrt, 2017-2018.
- [13] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(2):295–307, 2015.
- [14] Chao Dong, Chen Change Loy, and Xiaoou Tang. Accelerating the super-resolution convolutional neural network. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part II 14*, pages 391–407. Springer, 2016.
- [15] Zongcai Du, Jie Liu, Jie Tang, and Gangshan Wu. Anchor-based plain net for mobile image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2494–2502, 2021.
- [16] Dario Fuoli, Shuhang Gu, and Radu Timofte. Efficient video super-resolution through recurrent latent space propagation. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 3476–3485. IEEE, 2019.
- [17] Epic Games. Unreal engine 4.19: Screen percentage with temporal upsample, 2018.
- [18] GPUOpen-Effects. Gpuopen-effects/fidelityfx-fsr: Fidelityfx super resolution.
- [19] GPUOpen-Effects. Fidelityfx-fsr2/readme.md at master · gpuopen-effects/fidelityfx-fsr2, Oct 2022.
- [20] John K Haas. A history of the unity game engine. *Diss. Worcester Polytechnic Institute*, 483(2014):484, 2014.
- [21] Killian Herveau, Max Piochowiak, and Carsten Dachsbacher. Minimal convolutional neural networks for temporal anti aliasing. 2023.
- [22] Robert Herzog, Elmar Eisemann, Karol Myszkowski, and H-P Seidel. Spatio-temporal upsampling on the gpu. In *Proceedings of the 2010 ACM SIGGRAPH symposium on Interactive 3D Graphics and Games*, pages 91–98, 2010.
- [23] Jia-Bin Huang, Abhishek Singh, and Narendra Ahuja. Single image super-resolution from transformed self-exemplars. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5197–5206, 2015.
- [24] Yan Huang, Wei Wang, and Liang Wang. Bidirectional recurrent convolutional networks for multi-frame super-resolution. *Advances in neural information processing systems*, 28, 2015.
- [25] Zheng Hui, Xinbo Gao, Yunchu Yang, and Xiumei Wang. Lightweight image super-resolution with information multi-distillation network. In *Proceedings of the 27th acm international conference on multimedia*, pages 2024–2032, 2019.
- [26] Takashi Isobe, Xu Jia, Shuhang Gu, Songjiang Li, Shengjin Wang, and Qi Tian. Video super-resolution with recurrent structure-detail network. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XII 16*, pages 645–660. Springer, 2020.
- [27] Takashi Isobe, Fang Zhu, Xu Jia, and Shengjin Wang. Revisiting temporal modeling for video super-resolution. *arXiv preprint arXiv:2008.05765*, 2020.
- [28] Younghyun Jo, Seoung Wug Oh, Jaeyeon Kang, and Seon Joo Kim. Deep video super-resolution network using dynamic upsampling filters without explicit motion compensation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3224–3232, 2018.
- [29] Brian Karis. High-quality temporal supersampling. *Advances in Real-Time Rendering in Games, SIGGRAPH Courses*, 1(10.1145):2614028–2615455, 2014.

- [30] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [31] Wenbo Li, Xin Tao, Taian Guo, Lu Qi, Jiangbo Lu, and Jiaya Jia. Mucan: Multi-correspondence aggregation network for video super-resolution. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part X 16*, pages 335–351. Springer, 2020.
- [32] Jingyun Liang, Jiezhong Cao, Yuchen Fan, Kai Zhang, Rakesh Ranjan, Yawei Li, Radu Timofte, and Luc Van Gool. Vrt: A video restoration transformer. *arXiv preprint arXiv:2201.12288*, 2022.
- [33] Jingyun Liang, Yuchen Fan, Xiaoyu Xiang, Rakesh Ranjan, Eddy Ilg, Simon Green, Jiezhong Cao, Kai Zhang, Radu Timofte, and Luc Van Gool. Recurrent video restoration transformer with guided deformable attention. *arXiv preprint arXiv:2206.02146*, 2022.
- [34] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 136–144, 2017.
- [35] Anran Liu, Yihao Liu, Jinjin Gu, Yu Qiao, and Chao Dong. Blind image super-resolution: A survey and beyond. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [36] Edward Liu. Dlss 2.0-image reconstruction for real-time rendering with deep learning. In *GPU Technology Conference (GTC)*, 2020.
- [37] Jie Liu, Jie Tang, and Gangshan Wu. Residual feature distillation network for lightweight image super-resolution. In *Computer Vision–ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 41–55. Springer, 2020.
- [38] Zhi-Song Liu, Li-Wen Wang, Chu-Tak Li, and Wan-Chi Siu. Hierarchical back projection network for image super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 0–0, 2019.
- [39] Timothy Lottes. Fxaa. In *Nvidia White Paper*, 2009.
- [40] David Martin, Charless Fowlkes, Doron Tal, and Jitendra Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, volume 2, pages 416–423. IEEE, 2001.
- [41] Ben Mildenhall, Jonathan T Barron, Jiawen Chen, Dillon Sharlet, Ren Ng, and Robert Carroll. Burst denoising with kernel prediction networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2502–2510, 2018.
- [42] Markus Nagel, Marios Fournarakis, Rana Ali Amjad, Yelysei Bondarenko, Mart Van Baalen, and Tijmen Blankevoort. A white paper on neural network quantization. *arXiv preprint arXiv:2106.08295*, 2021.
- [43] 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*, June 2019.
- [44] Diego Nehab, Pedro V Sander, Jason Lawrence, Natalya Tatarchuk, and John R Isidoro. Accelerating real-time shading with reverse reprojection caching. In *Graphics hardware*, volume 41, pages 61–62, 2007.
- [45] Qualcomm. Efficient real-time int4 4k super-resolution on mobile. In *NeurIPS Demo Track*, 2022.
- [46] Alexander Reshetov. Morphological antialiasing. In *Proceedings of the Conference on High Performance Graphics 2009*, pages 109–116, 2009.
- [47] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.
- [48] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *arXiv:2104.07636*, 2021.
- [49] Mehdi SM Sajjadi, Raviteja Vemulapalli, and Matthew Brown. Frame-recurrent video super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6626–6634, 2018.
- [50] Daniel Scherzer, Stefan Jeschke, and Michael Wimmer. Pixel-correct shadow maps with temporal reprojection and shadow test confidence. In *Proceedings of the 18th Eurographics conference on Rendering Techniques*, pages 45–50, 2007.
- [51] Shuwei Shi, Jinjin Gu, Liangbin Xie, Xintao Wang, Yujiu Yang, and Chao Dong. Rethinking alignment in video super-resolution transformers. *arXiv preprint arXiv:2207.08494*, 2022.
- [52] Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1874–1883, 2016.
- [53] Bharat Singh, Tim K Marks, Michael Jones, Oncel Tuzel, and Ming Shao. A multi-stream bi-directional recurrent neural network for fine-grained action detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1961–1970, 2016.
- [54] 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 16*, pages 402–419. Springer, 2020.
- [55] Manu Mathew Thomas, Gabor Liptor, Christoph Peters, Sungye Kim, Karthik Vaidyanathan, and Angus G Forbes. Temporally stable real-time joint neural denoising and super-sampling. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 5(3):1–22, 2022.
- [56] Yapeng Tian, Yulun Zhang, Yun Fu, and Chenliang Xu. Tdan: Temporally-deformable alignment network for video

- super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3360–3369, 2020.
- [57] Xintao Wang, Kelvin CK Chan, Ke Yu, Chao Dong, and Chen Change Loy. Edvr: Video restoration with enhanced deformable convolutional networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.
- [58] Bartłomiej Wronski, Ignacio Garcia-Dorado, Manfred Ernst, Damien Kelly, Michael Krainin, Chia-Kai Liang, Marc Levoy, and Peyman Milanfar. Handheld multi-frame super-resolution. *ACM Transactions on Graphics (ToG)*, 38(4):1–18, 2019.
- [59] Lei Xiao, Salah Nouri, Matt Chapman, Alexander Fix, Douglas Lanman, and Anton Kaplanyan. Neural supersampling for real-time rendering. *ACM Transactions on Graphics (TOG)*, 39(4):142–1, 2020.
- [60] Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. Video enhancement with task-oriented flow. *International Journal of Computer Vision*, 127:1106–1125, 2019.
- [61] Lei Yang, Shiqiu Liu, and Marco Salvi. A survey of temporal antialiasing techniques. In *Computer graphics forum*, volume 39, pages 607–621. Wiley Online Library, 2020.
- [62] Lei Yang, Diego Nehab, Pedro V Sander, Pitchaya Sitthiamorn, Jason Lawrence, and Hugues Hoppe. Amortized supersampling. *ACM Transactions on Graphics (TOG)*, 28(5):1–12, 2009.
- [63] Roman Zeyde, Michael Elad, and Matan Protter. On single image scale-up using sparse-representations. In *Curves and Surfaces: 7th International Conference, Avignon, France, June 24-30, 2010, Revised Selected Papers 7*, pages 711–730. Springer, 2012.
- [64] Shangchen Zhou, Jiawei Zhang, Jinshan Pan, Haozhe Xie, Wangmeng Zuo, and Jimmy Ren. Spatio-temporal filter adaptive network for video deblurring. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2482–2491, 2019.