

FAST: A Framework to Accelerate Super-Resolution Processing on Compressed Videos

Zhengdong Zhang, Vivienne Sze
Massachusetts Institute of Technology
{zhangzd, sze}@mit.edu

Abstract

State-of-the-art super-resolution (SR) algorithms require significant computational resources to achieve real-time throughput (e.g., 60Mpixels/s for HD video). This paper introduces FAST (Free Adaptive Super-resolution via Transfer), a framework to accelerate any SR algorithm applied to compressed videos. FAST exploits the temporal correlation between adjacent frames such that SR is only applied to a subset of frames; SR pixels are then transferred to the other frames. The transferring process has negligible computation cost as it uses information already embedded in the compressed video (e.g., motion vectors and residual). Adaptive processing is used to retain accuracy when the temporal correlation is not present (e.g., occlusions). FAST accelerates state-of-the-art SR algorithms by up to $15\times$ with a visual quality loss of 0.2dB. FAST is an important step towards real-time SR algorithms for ultra-HD displays and energy constrained devices (e.g., phones and tablets).

1. Introduction

Video content can often have lower resolution than the display, either due to the fact that there is limited communication bandwidth, or that the resolution of displays rises faster than the resolution at which video content is captured. Thus, upsampling needs to be performed to match the resolution of the content to the display.

Today, most televisions perform upsampling using simple interpolation plus a sharpening filter [1]; these single frame upsamplers, such as bicubic, sinc, Lanczos, Catmull-Rom [2], Mitchell-Netravali [3] are based on simple splines to enable real-time throughput. *Super-resolution (SR)* can deliver higher visual quality than this by exploiting the non-local similarity of patches or by learning the mapping from the low-resolution to high-resolution from external datasets [4]. However, SR algorithms are computationally more expensive and slower than simple filtering. For in-

stance, state-of-the-art neural network based SR algorithms (e.g., SRCNN [5]) require powerful GPUs that consume hundreds of watts to achieve real-time performance [6, 7]. As a result, state-of-the-art SR algorithms are typically not used in televisions and portable screens as their high computational complexity limits their throughput and results in high power consumption (e.g., the power consumption of portable devices is limited to a few watts).

To bridge this gap, this paper proposes a technique called *Free Adaptive Super-resolution via Transfer (FAST)* to accelerate existing SR algorithms. Our approach is inspired by many concepts from the area of video coding research. In particular, numerous techniques have been developed in video coding research to identify correlations between pixels in videos. As a result, a compressed video is comprised of a compact description of the structure and correlation of the pixels in the video.

The main contribution of this paper is a framework that can leverage this structure information that is available for free in the compressed video to accelerate image processing, specifically super-resolution. FAST uses this free pre-computed information in two different ways:

1. It uses the motion vectors to exploit the inter-frame similarities between adjacent frames in a video sequence. As a result, it can transfer the SR result from one frame to the other such that SR is only applied to a subset of frames. Since the compressed video already provides the correlation information for free, the transfer step has negligible run time and computation cost compared with the original SR algorithms. Figure 1(a) illustrates the main steps of FAST. It should be noted that the transfer step can introduce visual degradation, as motion compensation in video coding is not perfect. In this work, we propose several visual quality metrics that are easy to compute, and then adaptively enable and disable the transfer on a block-by-block basis based on these metrics.
2. FAST uses the non-overlapping block structure embedded in the compressed video, rather than the over-

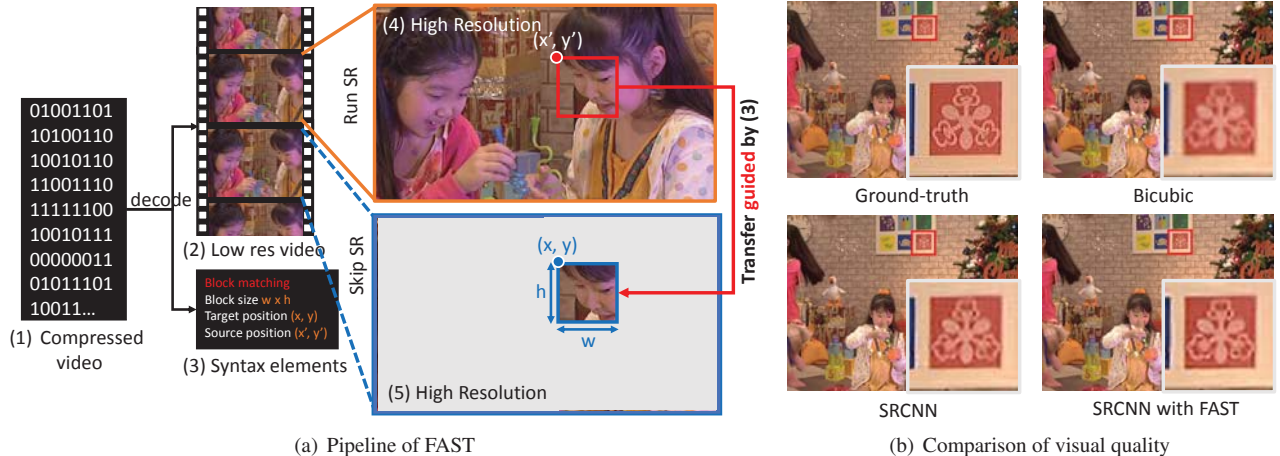


Figure 1. (a) **SR with FAST**: From the compressed video (1), the video decoder decompresses the low-resolution video frames (2) and the syntax elements (3). The SR algorithm generates the high resolution output for the first frame (4). Guided by the syntax elements (3), FAST adaptively transfers blocks from the first frame (4) to the second frame (5) rather than applying super-resolution on the second frame. (b) **Visual quality results**: Running SRCNN with FAST preserves the rich high frequency details that SRCNN generates compared with the blurry output of bicubic interpolation. Note that the complexity of FAST is comparable to bicubic interpolation on most of the transferred frames.

lapping blocks traditionally used for SR; using non-overlapping blocks accelerates the processing relative to the overlapping blocks. One of the drawbacks from non-overlapping blocks is that it can result in block artifacts. To address this, an adaptive deblocking filter is applied at the block edges to remove the block artifacts while retaining the true edges in the video.

We demonstrate the FAST framework acceleration using three well-known SR algorithms (KRR [8], ANR [9], SRCNN [5]) on the 20 video sequences that were used in the development of the latest video compression standard H.265/HEVC [10]; these sequences cover a wide range of videos including natural and screen content. On these sequences, FAST delivers up to $15\times$ acceleration with around 0.2dB drop in PSNR. Figure 1(b) shows how FAST maintains the visual quality of the super-resolution using SRCNN, which is significantly better than the bicubic upsampled result, even though the complexity of FAST is comparable to bicubic upsampling on most of the transferred frames. A demo video is available at [11].

2. Related Work on Super Resolution

There are two forms of super-resolution (SR) algorithms: single frame and multiple frame. Single-frame based SR algorithms are applied to each video frame independently, while multi-frame based SR algorithms are applied to multiple frames in a small time window. FAST can be applied to either SR approach.

Sophisticated single-frame based SR algorithms often leverage machine learning techniques [12, 13]. Among them are sparse-representation [14, 15], Kernel Ridge

Regression (KRR) [8], anchored neighbor regression (ANR) [9], and in-place example regression [16]. Recently, deep convolutional neural networks (CNNs) have been used to perform super-resolution (*e.g.*, SRCNN [5, 17]). They achieve state-of-the-art results at the cost of high computational complexity; for instance, the three convolutional layers in SRCNN [5] require 8032 multiplications per pixel, which is significantly higher than simple interpolation with one filter. While CNNs can be accelerated by powerful GPUs to achieve real-time performance on high-definition videos (1920×1080) [6, 7], these GPUs consume several hundred watts and are not suitable for embedding into televisions and portable devices (*e.g.*, phones, tablets).

There are also SR algorithms that exploit the self-similarities of blocks within each image [18, 19]. The proposed FAST framework shares many common insights with in-place example regression [16], which performs local block prediction inside the same frame; the important distinction is that FAST uses predictions across frames, which has negligible cost when exploiting embedded information in the compressed videos.

Multiple-frame based super-resolution algorithms are largely based on the registration of neighboring frames [20]. Many of these algorithms are iterative, including the Bayesian based approach [21], and the ℓ_1 -regularized total variation based approach [22]. At the same time, there are non-iterative methods that avoid registration with non-local mean [23], 3D steer kernel regression [24], and self-similarity [25]. Deep neural networks can also be used in the form of bidirectional recurrent convolutional networks [26], and deep draft-ensemble learning [27]. Spatial-

temporal coherence of multiple-frame based SR is addressed in [28] and [29]. Similar to this work, motion compensation is also used in [29] with two key differences: (1) both the original and motion compensated frames are fed into the neural network, while in FAST only the first frame is processed by the neural network; (2) the motion vectors are explicitly computed with optical flow, while FAST uses the motion vectors embedded in the compressed video. Thus, techniques proposed in FAST could potentially help accelerate [29].

3. Video Coding Basics

The FAST framework is inspired by many concepts from video coding research; in particular, the use of temporal correlation and the adaptive block-based processing. Temporal correlation is used in video coding during a process called motion compensation. Motion compensation is widely used in popular video coding standards including MPEG-2 [30], used for HDTV broadcast; H.264/AVC [31], used for most video content on the Internet; and H.265/HEVC [32], the most recent standard. Motion compensation involves predicting a block of pixels in the current frame from a block of pixels in a temporally-adjacent, previously-encoded frame.

During compression, videos are divided into non-overlapping blocks and encoded on a block-by-block basis. At the video encoder, motion estimation is used to determine which block of pixels in the previously encoded frame would best match the current block for minimum prediction error. As a result, the encoder only needs to signal the small difference between the two blocks, which typically requires fewer bits than signaling the original pixels; this difference is referred to as the *residual*. The offset between the location of the source block in the previous frame and the location of the target block in the current frame is called a *motion vector*. Both the residual and motion vector are part of the *syntax elements* that are embedded in the compressed video to express the structure of the video as shown in Figure 1(a). The video decoder uses these syntax elements to reconstruct the video; for instance, the motion compensation at the decoder uses the motion vector to determine which pixels to use from the previously encoded frame to generate the source block, and adds the residual to reconstruct the target block.

Video encoders divide up a video into sets of frames called group-of-pictures (GOP); motion compensation is then applied to frames within the same GOP. Each GOP typically contains between 6 to 16 frames that are visually similar with no scene transitions [33]. FAST will use the GOP size to define the maximum distance for transfer. The blocks of pixels in the first frame of a GOP are encoded by only exploiting spatial correlation (*i.e.*, correlation with neighboring pixels within the frame), which is referred to as intra-predicted blocks; these frames are referred to as

I frames. The subsequent frames are primarily composed of blocks that exploit temporal correlation, referred to as motion compensated blocks; these frames are referred to as P frames and may contain some intra-predicted blocks in cases of low temporal correlation.

4. FAST Framework

The FAST framework consists of the following features:

- Transfer SR pixels using motion compensation
- Adaptive transfer to retain visual quality (*i.e.*, PNSR)
- Apply SR to non-overlapping blocks and remove block artifacts with an adaptive deblocking filter

4.1. Transfer using Motion Compensation

To perform a transfer, the SR pixels are copied based on the scaled motion vector to the previous frame; here, the *scaled* motion vector refers to the fact that the motion vector extracted from the compressed video is multiplied by $2\times$ if FAST is performing a $2\times$ upsampling. If the scaled motion vector is fractional, an interpolation filter is applied to the SR pixels; FAST uses the same interpolation filter as the one used for fractional motion vectors during motion compensation for video decompression [34]. The residual is upsampled using bicubic interpolation and added to the transferred SR pixels.

Recall that the video encoder uses intra-prediction for blocks when there is not sufficient temporal correlation. In these cases, transfer using motion compensation cannot be used. However, with the exception of the first frame in the GOP (an I frame), intra-predicted blocks rarely occur in the subsequent P frames; this makes sense, as otherwise it would lead to poor coding efficiency of the video. Since there are few intra-predicted blocks, FAST upsamples them with bicubic upsampling.

In summary, FAST only applies the expensive SR processing on the first frame of a GOP (except when the transfer needs to be reset as discussed in Section 4.2.2). For all other frames in the GOP, FAST uses lightweight processing such as interpolation and bicubic upsampling. Since a typical GOP size is on the order of 6 to 16 frames [33], the majority of the frames are accelerated with lightweight processing; the overall acceleration depends on the GOP size. Section 5 reports the visual quality impact and acceleration of FAST for various GOP sizes and video sequences.

4.2. Adaptive Transfer

Transferring using motion compensation is not perfect. Error is introduced due to (1) poor temporal correlation where the motion compensated SR pixels are not good predictors; (2) the upsampling of the residual using bicubic interpolation does not account for potential high frequency

components in the prediction error. The error from adaptive transfer is more noticeable when the quantization used during compression is low since there are more high frequency components in the residual that are not preserved by the bicubic interpolation.

Figure 2(a) shows an example of when the error from the transfer, specifically the bicubic upsampling of the residual, can cause ringing artifacts. These types of errors occur when the source block has a sharp edge and the target block should be smooth; this mismatch between source and target can occur due to the rate-distortion optimization in the video encoder [35], where both the visual quality and the bits of the syntax elements must be considered (*e.g.*, trade-off fewer bits for the motion vector for more distortion and larger residual). The bicubic upsampling smooths out the sharp edge in the residual, which results in an incorrect sharp edge in reconstructed output as shown in Figure 2(b).

4.2.1 Deciding which blocks to transfer

In order to maintain visual quality, only a subset of the blocks in a frame are transferred. These blocks are adaptively selected based on the magnitude of the residual; note the value of the residual can be obtained for free from the video decoder. Specifically, transfer is disabled for blocks with large residuals (*i.e.*, *mean absolute value of their residual*); these disabled blocks are upsampled with bicubic interpolation. A threshold η is used to detect this blocks; the value of η is trained on the Middlebury stereo dataset [36] to maximize the PSNR across all blocks. In all the experiments presented in this paper, we set $\eta = 10$. Figure 2(a) shows how this simple threshold avoids the ringing artifact in the FAST result; this approach increases the PSNR by around 0.17dB. Unlike other approaches that use thresholding for acceleration [16, 37], FAST uses thresholding to improve visual quality.

4.2.2 Deciding when to reset transfer

While the residual error may be low for a given frame, it can accumulate as the number of transferred frames increase. This is due to the fact that the frames can be chained in a I-P-P-P GOP structure. Thus, the transfer should also be adaptively reset when the accumulated error gets too large. This accumulated error needs to be carefully modeled as it involves the interaction of error in different frames. As discussed earlier, the upsampling of the residual using bicubic interpolation does not account for potential high frequency components in the prediction error; accordingly, we use the Laplacian of the residual which has a higher correlation (0.484) with the PSNR drop compared with other metrics such as the magnitude of the residual (0.321). Thus the accumulated high frequency error can be modeled as the accumulated error of its source block from the previous frame

plus the Laplacian of the residual for the target block in the current frame. This computation only involves simple filtering with a Laplacian filter and an addition to accumulate the error. When this error exceeds a given threshold, the transfer is reset, and the accumulated error is reset to zero.

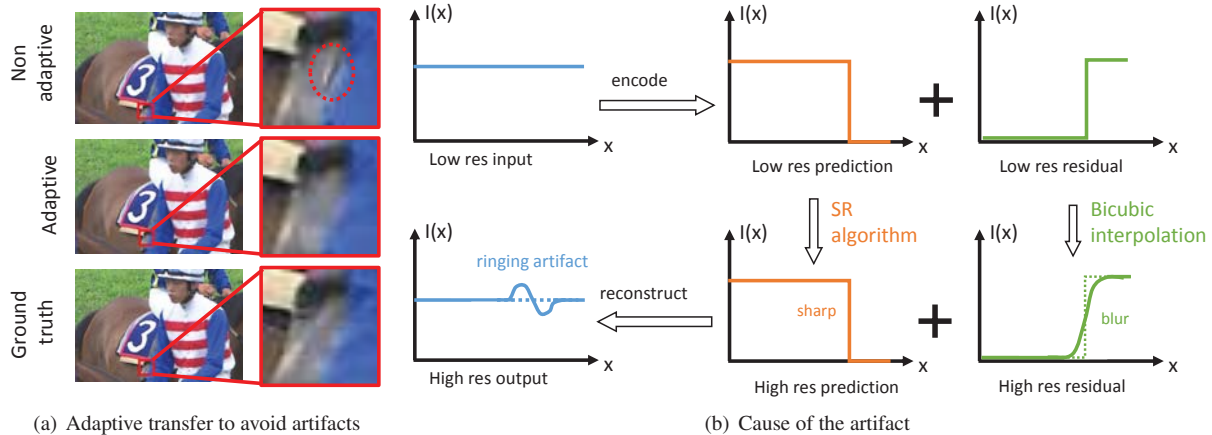
The reset involves applying super-resolution on the low resolution decoded block and thus has an impact on the overall throughput. Sweeping the reset threshold gives an accuracy versus speed trade-off which is discussed in Section 5 and shown in Figure 7.

4.3. Non-Overlapping Blocks

Most existing SR algorithms divide an image into densely overlapped blocks, and average the outputs of these overlapped blocks to avoid discontinuities at the block boundaries. This is computationally expensive since each pixel in a frame is processed multiple times since it belongs to multiple blocks [13, 14, 16, 37]. In contrast, video coding uses non-overlapping block, as shown in Figure 3(a), so that each pixel is covered by exactly one block. FAST uses this non-overlapping block structure, where each pixel is only processed once, which significantly reduces the computational complexity of the transferred frames.

The drawback of using non-overlapping blocks is that it can introduce artificial edges at the block boundaries. FAST addresses this by applying an adaptive deblocking filter, specifically the one used in H.265/HEVC [38], on the block edges; adaptive deblocking filters are widely used in video coding and have been shown to give similar visual quality [34] as using overlapped blocks [39]. The main objective of the deblocking filter is to remove the artificial edges on the block boundary due to non-overlapping block-based coding while keeping the true edges. Adaptive processing is used to decide whether to apply the deblocking filter to each block edge, and if so, decide the smoothing strength of the filter. This decision is based on the differences in coding mode and pixels on both sides of the block boundary including motion vectors, whether intra-prediction was used, the number of non-zero coefficients, the amount of quantization, and the variations of the pixels [38].

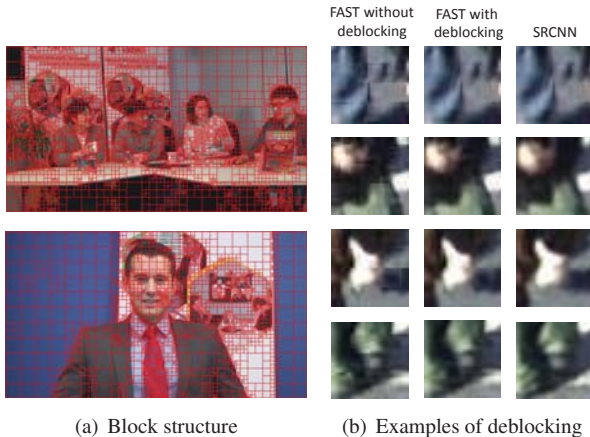
Using the H.265/HEVC deblocking filter to remove the artificial edges, as shown in Figure 3(b), increases the PSNR of the output by around 0.3dB, which is only 0.03dB less than using SRCNN with overlapping blocks. At the same time, using non-overlapping blocks can give over an order of magnitude speed up as compared using overlapping blocks (*e.g.*, each pixel in the first convolutional layer of SRCNN is processed 81 times). The deblocking filter is lightweight, and accounts for only 17% of the video decoder processing time on an ARM CPU [40], and only 3% of the video decoder power consumption in specialized hardware [41].



(a) Adaptive transfer to avoid artifacts

(b) Cause of the artifact

Figure 2. (a) Adaptive transfer avoids the ringing artifact in flat areas of the output of non-adaptive transfer. (b) Sharp edges in the residual causes ringing artifacts.



(a) Block structure

(b) Examples of deblocking

Figure 3. (a) An image is adaptively divided into non-overlapping blocks, with larger blocks corresponding to simple and well-predicted content. (b) Examples of running SRCNN with FAST before and after deblocking, compared to SRCNN results on the second frame. (*Best viewed in color*)

5. Experimental Results

5.1. Evaluation Dataset and Setup

We evaluate the FAST framework on the 20 video sequences used in the development of the latest video coding standard H.265/HEVC, as they cover a wide range of video content [10]. We use the original uncompressed videos as the high-resolution ground-truth. We then synthetically downsample them to lower-resolution, and encode them with the H.265/HEVC encoder [42]. The FAST framework uses the decompressed low resolution frames and the decoded syntax elements as the input, which can be obtained for free from the video decoder.

We evaluate the FAST framework on three SR algo-

gorithms: KRR [8], ANR [9] and SRCNN [5]¹. For each algorithm, we conduct two experiments: (1) we run the SR algorithm directly on all of the low-resolution frames; (2) we use the SR algorithm to upsample the first frame, then use FAST to transfer the result to all the other frames.

For quantitative evaluation, we compute the PSNR between each output frame and the ground-truth high-resolution frame. Since the ground-truth frames are not compressed, the reported PSNR includes both the quantization noise from lossy compression and the quality of the super-resolution algorithm. As a result, the reported gain of SR algorithms against bicubic interpolation in our experiments is around 1 dB lower than the gains reported in the original papers. We measure the run time to quantify the FAST acceleration using MATLAB on a 3.3GHz Xeon CPU.

5.2. Evaluation Results

Table 1 and Table 2 shows the PSNR and acceleration ratio for the various SR algorithms, when FAST transfers all the way to the 4th and 16th frame, respectively. Figure 4 shows the visual quality of the FAST framework on SRCNN for both the 2nd and the 16th frames. From these results we can see that the FAST framework is robust as it maintains the PSNR and achieves significant acceleration across all evaluated video sequences and SR algorithms.

Figure 5 shows that as the number of transferred frames increases (*i.e.*, GOP size), the acceleration increases, but so does the drop on in PSNR. This can be explained by looking at the breakdown in PSNR and acceleration per frame as shown in Table 3 and Table 4, respectively. In the first four frames, FAST gives comparable PSNR as SRCNN. At the 16th frame, FAST gives slightly worse PSNR (< 0.2 dB) than SRCNN, which is still significantly better than bicubic

¹Only the MATLAB code for SRCNN is publicly available; thus, our run times are higher than numbers reported in [5] which used C code.

Sequence name	Size	KRR			SRCNN			ANR			Bicubic PSNR
		SR PSNR	FAST PSNR	Speed up	SR PSNR	FAST PSNR	Speed up	SR PSNR	FAST PSNR	Speed up	
BQMall	416×240	28.67	28.63	3.9×	28.92	28.88	4.1×	28.5	28.48	4.1×	27.65
BQSquare	208×112	24.1	24.16	3.9×	24.78	24.85	4.1×	23.67	23.75	4.1×	22.77
BQTerrace	960×528	29.01	29.22	3.9×	29.37	29.51	4.0×	28.73	28.97	3.9×	27.45
BasketballDrill	416×240	31.7	31.69	3.9×	31.6	31.6	3.6×	31.18	31.18	3.8×	30.19
BasketballDrillText	416×240	31.7	31.69	3.9×	31.6	31.6	4.1×	31.18	31.18	3.8×	30.19
BasketballDrive	960×528	34.51	34.51	3.6×	34.65	34.65	3.6×	34.19	34.21	3.9×	33.49
BasketballPass	208×112	30.54	30.58	3.9×	30.63	30.65	3.7×	30.34	30.4	3.6×	29.7
BlowingBubbles	208×112	28.68	28.76	3.8×	28.68	28.74	3.7×	28.52	28.57	3.7×	27.91
Cactus	208×112	32.68	32.8	3.9×	32.72	32.83	3.5×	32.59	32.7	3.9×	31.78
ChinaSpeed	960×528	24.7	24.66	3.9×	24.9	24.83	4.0×	24.52	24.52	4.4×	23.97
FourPeople	512×384	34.43	34.43	4.1×	34.75	34.76	3.9×	34.36	34.36	3.9×	33
Johnny	640×352	37.08	37.09	4.0×	37.22	37.23	3.5×	36.81	36.84	4.1×	35.76
Kimono	960×528	37.5	37.44	3.7×	37.36	37.38	3.7×	37.44	37.4	3.9×	37.56
KristenAndSara	640×352	34.98	34.98	3.9×	35.84	35.85	4.3×	34.68	34.69	3.9×	33.22
ParkScene	960×528	33.35	33.36	3.9×	33.35	33.37	3.7×	33.33	33.36	4.0×	32.98
PartyScene	416×240	25.41	25.4	3.9×	25.58	25.58	4.3×	25.24	25.23	3.4×	24.55
PeopleOnStreet	1280×800	33.54	33.34	3.8×	33.35	33.19	4.0×	33.67	33.45	3.5×	32.89
RaceHorses	416×240	29.9	29.79	3.6×	29.95	29.82	3.9×	29.77	29.7	2.8×	29.02
SlideEditing	640×352	19.93	19.93	4.0×	20.73	20.73	3.5×	19.86	19.87	3.6×	19.3
Traffic	1280×800	34.73	34.71	3.9×	34.73	34.72	4.0×	34.75	34.74	3.9×	33.97
Average		30.86	30.86	3.9×	31.04	31.04	3.9×	30.67	30.68	3.8×	29.87

Table 1. With 3 transfers from the 1st frame to the 4th frame, FAST gets around 4× speed up uniformly across all sequences for all SR algorithms, with no visual quality loss.

Sequence name	Size	KRR			SRCNN			ANR			Bicubic PSNR
		SR PSNR	FAST PSNR	Speed up	SR PSNR	FAST PSNR	Speed up	SR PSNR	FAST PSNR	Speed up	
BQMall	416×240	28.63	28.4	14.1×	28.85	28.6	15.2×	28.47	28.28	13.2×	27.68
BQSquare	208×112	23.98	24.1	14.7×	24.65	24.61	13.7×	23.59	23.78	9.2×	22.7
BQTerrace	960×528	28.98	29.16	14.9×	29.34	29.36	15.9×	28.69	28.95	14.8×	27.43
BasketballDrill	416×240	31.53	31.28	13.3×	31.41	31.21	12.7×	31.03	30.87	11.2×	30.08
BasketballDrillText	416×240	31.53	31.28	13.3×	31.41	31.21	13.9×	31.03	30.87	11.5×	30.08
BasketballDrive	960×528	34.43	34.22	12.1×	34.56	34.3	15.3×	34.13	34.02	14.5×	33.47
BasketballPass	208×112	30.56	30.26	13.9×	30.62	30.25	13.3×	30.37	30.18	8.7×	29.81
BlowingBubbles	208×112	28.45	28.35	14×	28.47	28.34	12.3×	28.31	28.25	8×	27.75
Cactus	208×112	32.46	32.36	14.1×	32.52	32.39	13.9×	32.37	32.32	14.4×	31.6
ChinaSpeed	960×528	24.79	24.65	13.8×	24.97	24.81	13.7×	24.61	24.54	15×	24.05
FourPeople	512×384	34.33	34.23	15.1×	34.64	34.55	15.1×	34.26	34.17	15.2×	32.91
Johnny	640×352	36.92	36.76	15.1×	37.06	36.89	14.8×	36.62	36.51	16.1×	35.54
Kimono	960×528	37.08	36.59	13.5×	36.94	36.56	13.5×	37.02	36.57	14.2×	37.17
KristenAndSara	640×352	34.9	34.78	14×	35.74	35.6	17.6×	34.6	34.51	15.4×	33.15
ParkScene	960×528	33.08	32.83	13.3×	33.07	32.82	13.4×	33.06	32.84	14.8×	32.74
PartyScene	416×240	25.01	24.85	13.3×	25.17	24.95	14×	24.88	24.78	9.3×	24.28
PeopleOnStreet	1280×800	33.35	32.48	12.7×	33.15	32.38	15.3×	33.47	32.55	9.4×	32.74
RaceHorses	416×240	29.91	29.31	12×	29.95	29.31	10.7×	29.78	29.31	4.4×	29.1
SlideEditing	640×352	19.88	19.84	16×	20.7	20.61	13.3×	19.8	19.79	11.5×	19.24
Traffic	1280×800	34.52	34.25	14.1×	34.51	34.25	15.8×	34.51	34.27	12.8×	33.78
Average		30.72	30.5	13.9×	30.89	30.65	14.2×	30.53	30.37	12.2×	29.77

Table 2. With 15 transfers from the 1st frame to the 16th frame, FAST gets more than 10× speed up on average over all sequences for all SR algorithms, with around 0.2dB PSNR loss. Nevertheless, the PSNR of FAST output is still significantly higher than the bicubic output. Note that the complexity of FAST is comparable to bicubic interpolation on most of the transferred frames.




Figure 4. Running different SR algorithms with FAST on different frames of different sequences. Top: *Cactus*. Mid: *BasketballDrill*. Bottom: *BQTerrace*. Note how FAST maintains the appearance of SR output. (*Best viewed in color*)

even though the transferred frames in FAST have similar complexity to bicubic. The average run time on the first 4 and 16 frames is also included. We can see that the cost of the transfer is negligible compared with SRCNN. Therefore, the average run time per frame is approximately the run time of applying SR to the first frame divided by the number of processed frames.

Note that in our evaluation we are reporting the PSNR drop for the worst case, where all frames are chained together in a I-P-P-P GOP structure. For other more flexible GOP structures, such as those that allow bi-prediction, we would expect the PSNR to reduce more slowly than I-P-P-P.

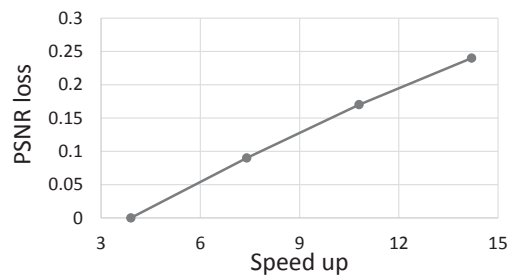


Figure 5. The higher the acceleration ratio, the more PSNR loss, albeit the PSNR loss is only around 0.2dB over all sequences even near 15 \times acceleration.

Frame ID	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	Mean 4	Mean 16
SRCNN	34.97	34.75	34.71	34.57	34.86	34.65	34.6	34.46	34.78	34.62	34.56	34.43	34.74	34.57	34.51	34.4	34.75	34.64
FAST	34.97	34.76	34.73	34.59	34.79	34.61	34.53	34.42	34.63	34.5	34.43	34.32	34.54	34.39	34.33	34.24	34.76	34.55
Bicubic	33.1	33.01	32.98	32.91	33.01	32.91	32.87	32.8	32.97	32.89	32.84	32.79	32.95	32.87	32.82	32.78	33	32.91

Table 3. PSNR (in dB) and visual quality loss of the FAST framework using SRCNN on the test sequence "FourPeople".

Frame ID	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	Mean 4	Mean 16
SRCNN	97.19	92.14	95.54	94.19	86.53	100.21	103.49	94.93	93.71	95.44	98.27	97.26	109.08	68.78	76.55	106.63	94.77	94.37
FAST	97.19	0.12	0.17	0.14	0.32	0.12	0.16	0.14	0.36	0.13	0.17	0.15	0.41	0.13	0.19	0.15	24.41	6.25
Speed up	1×	776×	556×	660×	268×	819×	663×	691×	259×	754×	595×	653×	266×	518×	399×	731×	3.9×	15.1×

Table 4. Run time (in msec) and acceleration of the FAST framework using SRCNN on the test sequence "FourPeople".

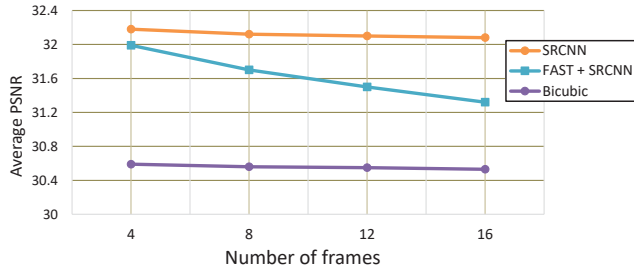
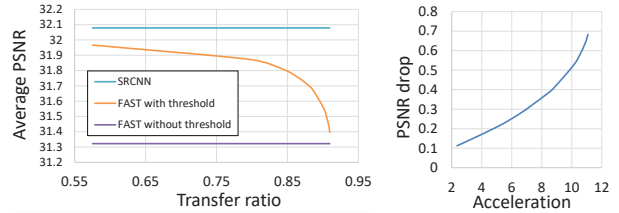


Figure 6. When $QP = 22$, the average PSNR across all the test sequences drops as FAST transfers more frames if no thresholding on accumulated error (*i.e.*, adaptive transfer reset) is performed.

The previous results were evaluated on video sequences with a quantization parameter (QP) setting of 27; this is a typical quantization setting used in video compression. To evaluate the sensitivity to the accumulated transfer error, we reduce the QP to 22, which reduces the compression ratio resulting in a higher quality video. Without adaptive transfer reset, we notice a large drop in PSNR as shown in Figure 6, where the 16th frame has around 0.7dB lower PSNR than SRCNN, which is much worse than the 0.2dB for QP of 27. As discussed earlier, this error accumulation is expected when the quantization used during compression is low since there are more high frequency components in the residual that are not preserved by the bicubic interpolation.

Applying the adaptive transfer reset reduces this PSNR drop at the cost of acceleration as shown in Figure 7. In these plots, we sweep the threshold discussed in Section 4.2.2. Decreasing the threshold reduces the number of blocks that can be transferred, thus improving the super-resolution visual quality while lowering the acceleration ratio. When more than 80% of the blocks are transferred for 5× acceleration, the PSNR drop is only around 0.2dB. Note that when the threshold is set to zero, there are still more than 50% blocks transferred, resulting in a 2× acceleration; this is due to the fact that on average, in a typical encoded video, over 50% of the pixels have zero residual; in all these cases, the blocks are transferred regardless of the threshold. Accordingly, the minimum transfer ratio plotted in Figure 7 is 0.55; the PSNR would remain the same below this ratio.

In summary, FAST accelerates SR algorithms by an order of magnitude in typical video quality settings. Even



(a) Average PSNR vs. transfer ratio (b) PSNR drop vs. acceleration ratio

Figure 7. Trade-off between visual quality and speed for FAST when $QP = 22$ for a GOP size of 16.

in the challenging setting with low quantization and the chained GOP structure (IPPP), FAST still achieves a significant (5×) acceleration with an acceptable PSNR loss around 0.2dB. This validates the effectiveness of FAST on accelerating a SR algorithms across all settings in practice.

6. Conclusions

This paper offers new insight on how to view the input to computer vision and image processing algorithms. Most visual content is stored in a compressed format; thus, rather than viewing the input as only an array of pixels, we should also take into account that it also contains useful information about the structure of the visual content.

In this paper, we have shown how exploiting this information in the FAST framework can accelerate various SR algorithms by up to 15× with only 0.2dB loss. Thus, the FAST framework is an important step towards enabling high visual quality SR algorithm in real-time for ultra-high resolution displays with low computational complexity.

FAST also demonstrates how using non-overlapping block division with deblocking filter reduces computation and avoids artifacts near the block boundary. This approach can potentially be extended to other SR and image processing algorithms.

As far as we know, FAST is the first technique to use the embedded information in compressed video to accelerate super-resolution on video. We believe that this approach can also be extended to accelerate other computer vision and image processing algorithms, which is becoming increasingly important as the latest state-of-the-art algorithms tend to use CNNs with high computational complexity.

References

- [1] J. Li, Y. Koudota, M. Barkowsky, H. Primon, and P. Le Callet, "Comparing upscaling algorithms from HD to Ultra HD by evaluating preference of experience," in *International Workshop on Quality of Multimedia Experience (QoMEX)*, 2014.
- [2] E. Catmull and R. Rom, "A Class of Local Interpolating Splines," *Computer Aided Geometric Design*, vol. 74, pp. 317–326, 1974.
- [3] D. P. Mitchell and A. N. Netravali, "Reconstruction filters in computer-graphics," *ACM Siggraph Computer Graphics*, vol. 22, pp. 221–228, 1988.
- [4] J. Yang and T. Huang, "Image super-resolution: Historical overview and future challenges," *Super-resolution imaging*, pp. 1–34, 2010.
- [5] C. Dong, C. C. Loy, K. He, and X. Tang, "Learning a deep convolutional network for image super-resolution," in *IEEE European Conference on Computer Vision (ECCV)*, 2014.
- [6] C. Dong, C. C. Loy, and X. Tang, "Accelerating the super-resolution convolutional neural network," in *IEEE European Conference on Computer Vision (ECCV)*, October 2016.
- [7] W. Shi, J. Caballero, F. Huszar, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang, "Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [8] K. I. Kim and Y. Kwon, "Single-image super-resolution using sparse regression and natural image prior," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 6, pp. 1127–1133, 2010.
- [9] R. Timofte, V. Smet, and L. Gool, "Anchored neighborhood regression for fast example-based super-resolution," in *IEEE International Conference on Computer Vision (ICCV)*, 2013.
- [10] F. Bossen, "Common test conditions and software reference configurations," *document JCTVC-H1100*, Feb. 2012.
- [11] "FAST Project Website." <http://www.rle.mit.edu/eems/fast/>.
- [12] W. T. Freeman, E. C. Pasztor, and O. T. Carmichael, "Learning low-level vision," *International Journal of Computer Vision*, vol. 40, no. 1, pp. 25–47, 2000.
- [13] W. T. Freeman, T. R. Jones, and E. C. Pasztor, "Example-based super-resolution," *IEEE Computer Graphics and Applications*, vol. 22, no. 2, pp. 56–65, 2002.
- [14] J. Yang, J. Wright, T. S. Huang, and Y. Ma, "Image super-resolution via sparse representation," *IEEE Transactions on Image Processing*, vol. 19, no. 11, pp. 2861–2873, 2010.
- [15] R. Zeyde, M. Elad, and M. Protter, "On single image scale-up using sparse-representations," in *Curves and Surfaces*, pp. 711–730, Springer, 2010.
- [16] J. Yang, Z. Lin, and S. Cohen, "Fast image super-resolution based on in-place example regression," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [17] J. Bruna, P. Sprechmann, and Y. LeCun, "Super-Resolution with Deep Convolutional Sufficient Statistics," *International Conference on Learning Representations (ICLR)*, 2016.
- [18] D. Glasner, S. Bagon, and M. Irani, "Super-resolution from a single image," in *IEEE International Conference on Computer Vision (ICCV)*, 2009.
- [19] J.-B. Huang, A. Singh, and N. Ahuja, "Single image super-resolution from transformed self-exemplars," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5197–5206, 2015.
- [20] S. Baker and T. Kanade, *Super-Resolution Optical Flow*. Carnegie Mellon University, The Robotics Institute, 1999.
- [21] C. Liu and D. Sun, "A Bayesian approach to adaptive video super resolution," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 209–216, 2011.
- [22] D. Mitzel, T. Pock, T. Schoenemann, and D. Cremers, "Video Super Resolution Using Duality Based TV-L1 Optical Flow," in *Proceedings of the 31st DAGM Symposium on Pattern Recognition*, pp. 432–441, 2009.
- [23] M. Protter, M. Elad, H. Takeda, and P. Milanfar, "Generalizing the nonlocal-means to super-resolution reconstruction," *IEEE Transactions on Image Processing*, vol. 18, no. 1, pp. 36–51, 2009.
- [24] H. Takeda, P. Milanfar, M. Protter, and M. Elad, "Super-resolution without explicit subpixel motion estimation," *IEEE Transactions on Image Processing*, vol. 18, no. 9, pp. 1958–1975, 2009.

- [25] S. Jeong, I. Yoon, and J. Paik, "Multi-frame example-based super-resolution using locally directional self-similarity," *IEEE Transactions on Consumer Electronics*, vol. 61, no. 3, pp. 353–358, 2015.
- [26] Y. Huang, W. Wang, and L. Wang, "Bidirectional Recurrent Convolutional Networks for Multi-Frame Super-Resolution," in *Conference on Neural Information Processing Systems (NIPS)*, 2015.
- [27] R. Liao, X. Tao, R. Li, Z. Ma, and J. Jia, "Video Super-Resolution via Deep Draft-Ensemble Learning," in *IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [28] C.-C. Hsu, L.-W. Kang, and C.-W. Lin, "Temporally coherent superresolution of textured video via dynamic texture synthesis," *IEEE Transactions on Image Processing*, vol. 24, no. 3, pp. 919–931, 2015.
- [29] J. Caballero, C. Ledig, A. Aitken, A. Acosta, J. Totz, Z. Wang, and W. Shi, "Real-time video super-resolution with spatio-temporal networks and motion compensation," *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [30] D. Le Gall, "MPEG: A video compression standard for multimedia applications," *Communications of the ACM*, vol. 34, no. 4, pp. 46–58, 1991.
- [31] "Recommendation ITU-T H.264: advanced video coding for generic audiovisual services," tech. rep., ITU-T, 2003.
- [32] *High efficiency video coding*. ITU-T Recommendation H.265 and ISO/IEC 230082, April 2013.
- [33] "Final Cut Pro 7 User Manual." <https://documentation.apple.com/en/finalcutpro/usermanual/index.html>.
- [34] V. Sze, M. Budagavi, and G. J. Sullivan, eds., *High Efficiency Video Coding (HEVC): Algorithms and Architectures*. Integrated Circuits and Systems, Springer, 2014.
- [35] G. J. Sullivan and T. Wiegand, "Rate-distortion optimization for video compression," *IEEE Signal Processing Magazine*, vol. 15, no. 6, pp. 74–90, 1998.
- [36] "Middlebury Stereo Datasets." <http://vision.middlebury.edu/stereo/data/>.
- [37] J. Yang, Z. Wang, Z. Lin, S. Cohen, and T. Huang, "Coupled dictionary training for image super-resolution," *IEEE Transactions on Image Processing*, vol. 21, no. 8, pp. 3467–3478, 2012.
- [38] A. Norikin, G. Bjontegaard, A. Fuldseth, M. Narroschke, M. Ikeda, K. Andersson, M. Zhou, and G. Van der Auwera, "Hevc deblocking filter," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 22, no. 12, pp. 1746–1754, 2012.
- [39] H. S. Malvar, "Lapped transforms for efficient transform/subband coding," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 38, no. 6, pp. 969–978, 1990.
- [40] F. Bossen, B. Bross, K. Suhring, and D. Flynn, "HEVC complexity and implementation analysis," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 22, no. 12, pp. 1685–1696, 2012.
- [41] M. Tikekar, C.-T. Huang, C. Juvekar, V. Sze, and A. P. Chandrakasan, "A 249-Mpixel/s HEVC video-decoder chip for 4K ultra-HD applications," *IEEE Journal of Solid-State Circuits*, vol. 49, no. 1, pp. 61–72, 2014.
- [42] G. J. Sullivan, J. Ohm, T. K. Tan, and T. Wiegand, "Overview of the high efficiency video coding (HEVC) standard," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 22, pp. 1649–1668, Dec 2012.