

AIM 2022 Challenge on Super-Resolution of Compressed Image and Video: Dataset, Methods and Results

Ren Yang, Radu Timofte, Xin Li, Qi Zhang, Lin Zhang, Fanglong Liu,
Dongliang He, Fu li, He Zheng, Weihang Yuan, Pavel Ostyakov, Dmitry Vyal,
Magauiya Zhussip, Xueyi Zou, Youliang Yan, Lei Li, Jingzhu Tang,
Ming Chen, Shijie Zhao, Yu Zhu, Xiaoran Qin, Chenghua Li, Cong Leng,
Jian Cheng, Claudio Rota, Marco Buzzelli, Simone Bianco,
Raimondo Schettini, Dafeng Zhang, Feiyu Huang, Shizhuo Liu,
Xiaobing Wang, Zhezhu Jin, Bingchen Li, Xin Li, Mingxi Li, Ding Liu,
Wenbin Zou, Peijie Dong, Tian Ye, Yunchen Zhang, Ming Tan, Xin Niu,
Mustafa Ayazoglu, Marcos Conde, Ui-Jin Choi, Zhuang Jia, Tianyu Xu,
Yijian Zhang, Mao Ye, Dengyan Luo, Xiaofeng Pan, and Liuhan Peng

Organizers: ren.yang@vision.ee.ethz.ch, radu.timofte@uni-wuerzburg.de

Abstract. This paper reviews the Challenge on Super-Resolution of Compressed Image and Video at AIM 2022. This challenge includes two tracks. Track 1 aims at the super-resolution of compressed image, and Track 2 targets the super-resolution of compressed video. In Track 1, we use the popular dataset DIV2K as the training, validation and test sets. In Track 2, we propose the LDV 3.0 dataset, which contains 365 videos, including the LDV 2.0 dataset (335 videos) and 30 additional videos. In this challenge, there are 12 teams and 2 teams that submitted the final results to Track 1 and Track 2, respectively. The proposed methods and solutions gauge the state-of-the-art of super-resolution on compressed image and video. The proposed LDV 3.0 dataset is available at https://github.com/RenYang-home/LDV_dataset. The homepage of this challenge is at https://github.com/RenYang-home/AIM22_CompressSR.

Keywords: Super-resolution, image compression, video compression

1 Introduction

Compression plays an important role on the efficient transmission of images and videos through the band-limited Internet. However, image and video compression unavoidably leads to compression artifacts, which may severely degrade the visual quality. Therefore, quality enhancement of compressed image and video has become a popular research topic. However, in the early years, due to the limitation of devices and band-width, the image and video are usually with low resolution. Therefore, when we intend to restore them to high resolution and good quality, we face the challenge to achieve both super-resolution and quality enhancement of compressed image (Track 1) and video (Track 2).

In the past decade, a great number of works were proposed for single image super-resolution [23,40,60,82,83,84,21,48] and there are also plenty of methods proposed for the reduction of JPEG artifacts [22,82,61,25,38]. Recently, the blind super-resolution [28,69,81] methods have been proposed. They are able to use one model to jointly handle the tasks of super-resolution, deblurring, JPEG artifacts reduction, etc. Meanwhile, video super-resolution [8,62,65,37,9,10,47,50] and compression artifacts reduction [77,76,78,30,71,68,70] also has become a popular topic, which aims at adequately exploring the temporal correlation among frames to facilitate the super-resolution and quality enhancement of videos. NTIRE 2022 [75] is the first challenge we organized on super-resolution of compressed video. The winner method [85] in the NTIRE 2022 challenge successfully outperforms the state-of-the-art method [11].

The AIM 2022 Challenge on Super-Resolution of Compressed Image and Video is one of the AIM 2022 associated challenges: reversed ISP [18], efficient learned ISP [36], super-resolution of compressed image and video [74], efficient image super-resolution [32], efficient video super-resolution [33], efficient Bokeh effect rendering [34], efficient monocular depth estimation [35], Instagram filter removal [42].

The AIM 2022 Challenge on Super-Resolution of Compressed Image and Video steps forward for establishing a benchmark of the super-resolution of JPEG image (Track 1) and HEVC video (Track 2). The methods proposed in this challenge are also have the potential to solve various super-resolution tasks. In this challenge, Track 1 utilizes the DIV2K [1] dataset, and Track 2 uses the proposed LDV 3.0 dataset, which contains 365 videos with diverse content, motion, frame-rate, etc. In the following, we first describe the AIM 2022 Challenge, including the DIV2K [1] dataset and the proposed LDV 3.0 dataset. Then, we introduce the proposed methods and the results.

2 AIM 2022 Challenge

The objectives of the AIM 2022 challenge on Super-Resolution of Compressed Image and Video are: (i) to advance the state-of-the-art in super-resolution of compressed inputs; (ii) to compare different solutions; (iii) to promote the proposed LDV 3.0 dataset.

2.1 DIV2K [1] dataset

The DIV2K [1] dataset consists of 1,000 high-resolution images with diverse contents. In Track 1 of AIM 2022 Challenge, we use the training (800 images), validation (100 images) and test (100 images) sets of DIV2K for training, validation and test, respectively.

2.2 LDV 3.0 dataset

The proposed LDV 3.0 dataset is an extension of the LDV 2.0 dataset [75,73,72] with 30 additional videos. Therefore, there are totally 365 videos in the LDV 3.0

dataset. The same as LDV and LDV 2.0, the additional videos in LDV 3.0 are collected from YouTube [27], containing 10 categories of scenes, i.e., *animal*, *city*, *close-up*, *fashion*, *human*, *indoor*, *park*, *scenery*, *sports* and *vehicle*, and they are with diverse frame-rates from 24 fps to 60 fps. To ensure the high quality of the groundtruth videos, we only collect the videos with 4K resolution, and without obvious compression artifacts. We downscale the videos to further remove the artifacts, and crop the width and height of each video to the multiples of 8, due to the requirement of the HEVC test model (HM). Besides, we convert videos to the format of YUV 4:2:0, which is the most commonly used format in the existing literature. Note that all source videos in our LDV 3.0 dataset have the licence of *Creative Commons Attribution licence (reuse allowed)*¹, and our LDV 3.0 dataset is used for academic and research proposes.

The Track 2 of AIM 2022 Challenge has the same task as the Track 3 of our NTIRE 2022 Challenge [75]. Therefore, we use the training, validation and test sets of the Track 3 in NTIRE 2022 as the training set (totally 270 videos) for the Track 2 in AIM 2022 Challenge. All videos in the proposed LDV, LDV 2.0 and LDV 3.0 datasets and the splits in NTIRE 2021, NTIRE 2022 and AIM 2022 Challenges are publicly available at https://github.com/RenYang-home/LDV_dataset.

2.3 Track 1 – super-resolution of compressed image

JPEG is the most commonly used image compression standard. Track 1 targets the $\times 4$ super-resolution of the images compressed with JPEG with the quality factor of 10. Specifically, we use the following Python codes to produce the low resolution samples:

```
from PIL import Image
img = Image.open(path_gt + str(i).zfill(4) + '.png')
w, h = img.size
assert w % 4 == 0
assert h % 4 == 0
img = img.resize((int(w/4), int(h/4)), resample=Image.BICUBIC)
img.save(path + str(i).zfill(4) + '.jpg', "JPEG", quality=10)
```

In this challenge, we use the version 7.2.0 of the Pillow library.

2.4 Track 2 – super-resolution of compressed video

Track 2 has the same task as the Track 3 in NTIRE 2022 [75], which requires the participants to enhance and meanwhile $\times 4$ super-resolve the HEVC compressed video. In this track, the input videos are first downsampled by the following command:

¹ <https://support.google.com/youtube/answer/2797468?hl=en>

```
ffmpeg -pix_fmt yuv420p -s WxH -i x.yuv
-vf scale=(W/4)x(H/4):flags=bicubic x.down.yuv
```

where x , W and H indicates the video name, width and height, respectively. Then, the downsampled video is compressed by HM 16.20² at QP = 37 with the default Low-Delay P (LDP) setting (*encoder_lowdelay_P_main.cfg*). Note that, we first crop the groundtruth videos to make sure that the downsampled width ($W/4$) and height ($H/4$) are integer numbers.

3 Challenge results

3.1 Track 1

The PSNR results and the running time of Track 1 are shown in Table 1. In this track, we use the images that are directly upscaled by the bicubic algorithm as the baseline. As we can see from Table 1, all methods proposed in this challenge achieves > 1 dB PSNR improvement over the baseline. The PSNR improvement of the top 3 methods are higher than 1.3 dB over the baseline. The VUE Team achieves the best result, that is ~ 0.1 dB higher the runner-up method. We can also see from Table 1 that the top methods consume high time complexity, while the method of the Giantpandacv Team is the most time-efficient one, whose running time is significantly lower than the methods with higher PSNR. Note that, the data in Table 1 are provided by the participants, so the data may be obtained under different hardware and conditions. Therefore, Table 1 is only for reference. It is hard to guarantee the fairness in comparing time efficiency.

The test and training details are presented in Table 2. As Table 2 shows, most methods use extra training data to improve the performance. In this challenge, Flickr2K [63] is the most popular dataset used in training, in addition to the official training data provided by the organizers. In inference, the self-ensemble strategy [64] is widely utilized. It has been proved to be an effective skill to boost the performance of super-resolution.

3.2 Track 2

Table 3 shows the results of Track 2. Similar to Track 1, we use the videos that are directly upscaled by the bicubic algorithm as the baseline performance in this track. The winner team NoahTerminalCV improves the PSNR by more than 2 dB over the baseline, and it successfully beats the winner method [85] in NTIRE 2022, which can be seen as the state-of-the-art method. The IVL method has the very fast running speed, and it is able to achieve the real-time super-resolution on the test videos. In Table 4, we can see that the NoahTerminalCV Team uses a large training set, including 90,000 videos collected from YouTube [27]. This may be obviously beneficial for their test performance. Note that, the data in Table 4 are provided by the participants. It is hard to guarantee the fairness in comparing time efficiency.

² https://hevc.hhi.fraunhofer.de/svn/svn_HEVCSoftware/tags/HM-16.20

Table 1. Results of Track 1 ($\times 4$ super-resolution of JPEG image). The test input is available at <https://codalab.lisn.upsaclay.fr/competitions/5076>, and the researchers can submit their results to the “testing” phase at the CodaLab server to get the performance of their methods to compare with the numbers in this table.

| Team | PSNR (dB) | Running time (s) | Hardware |
|--------------------|-----------|------------------|------------------|
| VUE | 23.6677 | 120 | Tesla V100 |
| BSR [45] | 23.5731 | 63.96 | Tesla A100 |
| CASIA LCVG [58] | 23.5597 | 78.09 | Tesla A100 |
| SRC-B | 23.5307 | 18.61 | GeForce RTX 3090 |
| USTC-IR [44] | 23.5085 | 19.2 | GeForce 2080ti |
| MSDRSR | 23.4545 | 7.94 | Tesla V100 |
| Giantpandacv | 23.4249 | 0.248 | GeForce RTX 3090 |
| Aselsan Research | 23.4239 | 1.5 | GeForce RTX 2080 |
| SRMUI [17] | 23.4033 | 9.39 | Tesla A100 |
| MVideo | 23.3250 | 1.7 | GeForce RTX 3090 |
| UESTC+XJU CV | 23.2911 | 3.0 | GeForce RTX 3090 |
| cvlab | 23.2828 | 6.0 | GeForce 1080 Ti |
| Bicubic $\times 4$ | 22.2420 | - | - |

Table 2. Test and training details of Track 1 ($\times 4$ super-resolution of JPEG image).

| Team | Ensemble for test | Extra training data |
|--------------------|--|--|
| VUE | Flip/rotation | ImageNet [19], Flickr2K [63] |
| BSR | Flip/rotation, three models for voting | Flickr2K [63], Flickr2K-L ³ |
| CASIA LCVG | Flip/rotation, three models | ImageNet [19] |
| SRC-B | Flip/rotation | Flickr2K [63] |
| USTC-IR | Flip/rotation | Flickr2K [63] and CLIC datasets |
| MSDRSR | Flip/rotation | Flickr2K [63] and DIV8K [29] |
| Giantpandacv | Flip/rotation, TLC [15] | Flickr2K [63] |
| Aselsan Research | Flip/rotation | Flickr2K [63] |
| SRMUI | Flip/rotation | Flickr2K [63] and MIT 5K [7] |
| MVideo | Flip/rotation | - |
| UESTC+XJU CV | Flip/rotation | - |
| cvlab | - | 4,600 images |
| Bicubic $\times 4$ | - | - |

³Flickr2K-L is available as “flickr2k-L.csv” at https://github.com/RenYang-home/AIM22_CompressSR/

Table 3. Results of Track 2 ($\times 4$ super-resolution of HEVC video). **Blue** indicates the state-of-the-art method. The test input and groundtruth are available on the homepage (see the abstract) of the challenge.

| Team | PSNR (dB) | Time (s) | Hardware |
|-----------------------------|-----------|----------|------------------|
| NoahTerminalCV | 25.1723 | 10 | Tesla V100 |
| NTIRE'22 Winner [85] | 24.1097 | 13.0 | Tesla V100 |
| IVL | 23.0892 | 0.008 | GeForce GTX 1080 |
| Bicubic $\times 4$ | 22.7926 | - | - |

4 Teams and methods

4.1 VUE Team

The method proposed by the VUE Team is called TCIR: A Transformer and CNN Hybrid Network for Image Restoration. The architecture of TCIR is shown

Table 4. Test and training details of Track 2 ($\times 4$ super-resolution of HEVC video). Blue indicates the state-of-the-art method.

| Team | Ensemble for test | Extra training data |
|--|--|---|
| NoahTerminalCV NTIRE'22 Winner [85] | Flip/rotation Flip/rotation, two models | 90K videos ⁴ from YouTube [27] 870 videos from YouTube [27] |
| IVL | - | - |
| Bicubic $\times 4$ | - | - |

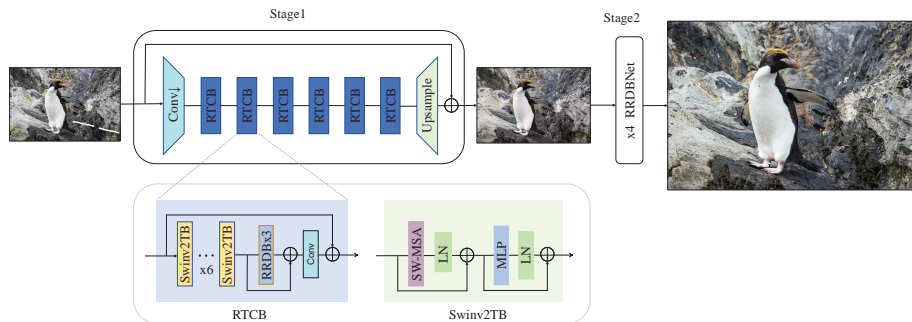
⁴Dataset is available as “dataset_Noah.txt” at https://github.com/RenYang-home/AIM22_CompressSR/

in Fig. 1. Specifically, they decouple this task of Track 1 into two sub-stages. In the first stage, they propose a Transformer and CNN hybrid network (TCIR) to remove JPEG artifacts, and in the second stage they use a finetuned RRDBNet for $\times 4$ super-resolution. The proposed TCIR is based on SwinIR [48] and the main improvements are as follows:

- 1) They conduct $\times 2$ downsampling to the JPEG-compressed input by a convolution with the stride of 2. The main purpose of this downsampling is to save GPU memory and speed up the model. Since the images compressed by JPEG with the quality factor of 10 are very blurry, this does not affect the performance of TCIR.
- 2) Then, they use the new Swinv2 transformer block layer to replace the STL in the original SwinIR to greatly improve the capability of the network.
- 3) In addition, they add several RRDB [84] modules to the basic blocks of TCIR and this combines the advantages of Transformer and CNN.

4.2 NoahTerminalCV Team

The method proposed by the NoahTerminalCV Team is called Enhanced Video Super-Resolution through Reference-Based Frame Refinement. As Fig. 2 shows, the proposed method consists of two subsequent stages. Firstly, they perform an

**Fig. 1.** Overview of the TCIR method proposed by the VUE Team.

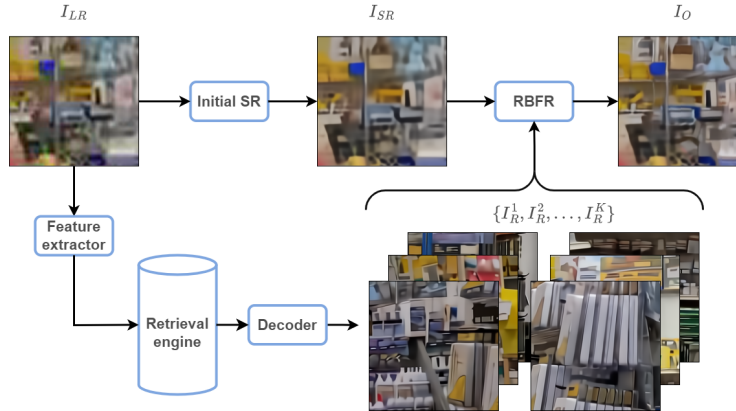


Fig. 2. The method proposed by the NoahTerminalCV Team.

initial super-resolution using a feed-forward multi-frame neural network. Then, the second step is called reference-based frame refinement. They find top K similar images for each low-resolution input frame from the external database. Then, they run a matching correction step for every patch on this input frame to perform a global alignment of reference patches. As a result, the I_{SR} , which comes from the first stage, and a set of globally aligned references $\{I_R^1, I_R^2, \dots, I_R^K\}$ are obtained. Finally, they are processed with RBF network (N_{RBF}) to handle residual misalignments and to properly transfer texture, details from reference images to initially super-resolved output I_{SR} . The details of training and test are described in the following.

4.2.1 Training

Initial Super-Resolution (Initial SR). The NoahTerminalCV Team upgraded the BasicVSR++ [10] by increasing channels to 128 and reconstruction blocks to 45. The BasicVSR++ is trained from scratch (except SPyNet) using a pixel-wise L_1 objective on the full input images without cropping and fine-tuned using $L_1 + L_2$ objectives. The training phase took about 21 days using 8 NVIDIA Tesla V100 GPUs. They observed a slight performance boost if the model is fine-tuned with a combination of L_1 and L_2 losses.

Reference-based Frame Refinement (RBFR). The information from subsequent frames is not always enough to produce a high-quality super-resolved image. Therefore, after initial super-resolution using upgraded BasicVSR++, they employed a reference-based refinement strategy. The idea is to design a retrieval engine that will find top K closest features in the database and then transfer details/texture from them to the initially upscaled frame. The retrieval engine includes feature extractor, database of features, and autoencoder.

Feature Extractor. They trained a feature extractor network that takes a low-resolution image I_{LR} and represents it as a feature vector. They used a

contrastive learning [12] framework to train the feature extractor: for positive samples, they used two random frames from the same video, while for the negative samples we employ frames from different videos. The backbone for a feature extractor was Resnet-34 [31].

Database and Autoencoder. The database consists of 2,000,000 samples generated from the training dataset. Each sample is compressed into a feature $z = N_E(I_{HR})$ using the Encoder N_E , since saving as images in a naive way is not practically plausible. Once they find top K similar features, the Decoder N_D is used to reconstruct the original input $\hat{I}_{HR} = N_D(z)$.

Retrieval Engine. After compressing the database of images using the trained Encoder, obtaining latent representations, and representing all low-resolution versions as a feature vector of size 100 extracted from the trained Feature Extractor, they build an index using the HNSW [54] algorithm from the nmslib [5] library. This algorithm allows searching for the top K nearest neighbors in the database.

RBFR. Finally, we train a network N_{RBFR} that takes the result of initial super-resolution I_{SR} and top K similar images from the database $\{I_R^1, I_R^2, \dots, I_R^K\}$. The network produces the final prediction I_O . We train N_{RBFR} through the L1 objective between I_O and I_{HR} . As a N_{RBFR} , we use NoahBurstSRNet [3] architecture, since it effectively handles small misalignments and can properly transfer information from reference non-aligned images.

4.2.2 Test

Initial Super-Resolution (Initial SR). During the inference, in order to upscale the key-frame I_{LR}^i , they put it to the initial super-resolution network together with additional frames $I_{LR}^{i-1}, I_{LR}^{i+1}, I_{LR}^{i-2}, I_{LR}^{i+2} \dots$. The number of additional frames during the inference is set to the full sequence size (up to 600 images).

Reference-based Frame Refinement (RBFR). For RBFR, top K (typically 16) similar images are first obtained using the retrieval engine. Then, the inference is done in a patch-wise manner. They extract a patch from the I_{SR} and use the Template Matching [6] to perform a global alignment and find the most similar patches on the images $\{I_R^1, I_R^2, \dots, I_R^K\}$. Then, they put them to the N_{RBFR} to generate the final result.

4.3 BSR Team [45]

For most low-level tasks, like image super-resolution, the network is trained on cropped patches rather than the full image. That means the network can only look at the pixels inside the patches during the training phase, even though the network’s ability is becoming more and more powerful and the receptive field of the deep neural network could be very large.

The patch size heavily affects the ability of the network. However, with the limitation of the memory and the computing power of GPU, it is not a sanity choose to train the network on the full image. To address the above-mentioned

problem, the BSR Team proposes a multi-patches method to greatly increase the receptive field in the training phase while increasing very little memory.

As shown in Fig. 3, they crop low-resolution input patch and its eight surrounding patches as our multi-patches network’s input. They use HAT [13] as the backbone, and propose Multi Patches Hybrid Attention Transformer (MPHAT). Compared with HAT [13], MPHAT just simply changew the input channel of the network for the multi-patch input. On the validation set of the challenge, the proposed MPHAT achieves the PSNR performance of 23.6266 dB, which is obviously higher than the HAT without multi-patches (23.2854 dB).

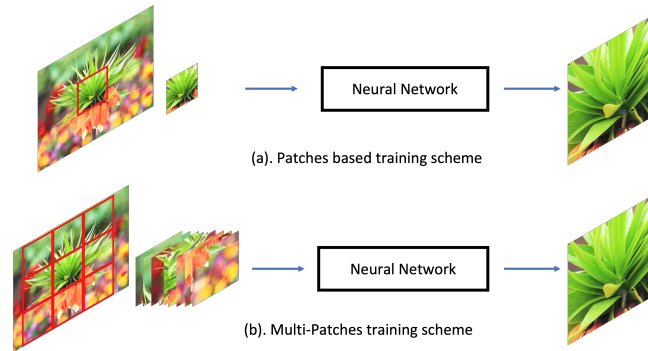


Fig. 3. Illustration of the multi patches scheme proposed by the BSR Team. Top images represent general patches based training scheme and bottom image represent multi-patches scheme. Low resolution input patch and its eight surrounding patches are cropped then send to the neural network to reconstruct the super resolution image of the centre patch. The neural network chosen in this competition is HAT [13].

In the training phase, they train the network by using Adam optimizer with $\beta_1 = 0.9$ and $\beta_2 = 0.99$ to minimize the MSE loss. The model is trained for 800,000 iterations with mini-batches of size 32 and patch size 64. The learning rate is initialized as 10^{-4} and reduced to half at the 300,000th, 500,000th, 650,000th, 700,000th, 750,000th iterations, respectively.

4.4 CASIA LCVG Team [58]

The CASIA LCVG Team proposes a consecutively-interactive dual-branch network (CIDBNet) to take advantage of both convolution and transformer operations, which are good at extracting local features and global interactions, respectively. To better aggregate the two-branch information, they newly introduce an adaptive cross-branch fusion module (ACFM), which adopts a cross-attention scheme to enhance the two-branch features and then fuses them weighted by a content-adaptive map. Experimental results demonstrate the effectiveness of CIDBNet, and in particular, CIDBNet achieves higher performance than a larger

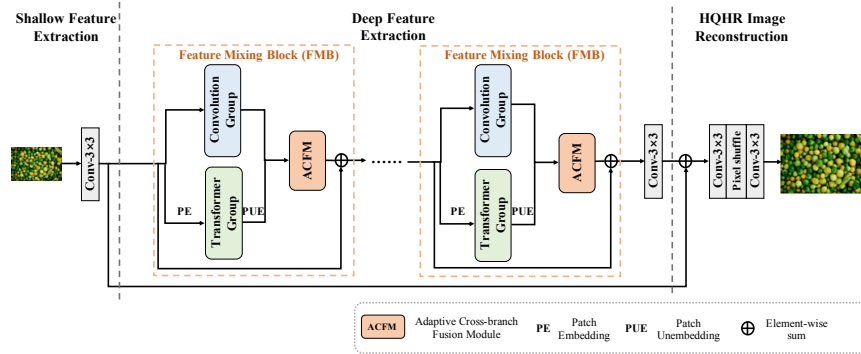


Fig. 4. The overall architecture of the Consecutively-Interactive Dual-Branch network (CIBDNet) of the CASIA LCVG Team.

variant of HAT (HAT-L) [13]. The framework of the proposed method is illustrated in Fig. 4.

They adopt 1,280,000 images from ImageNet [19] as training set and all the models are trained from scratch. They set the input patch size to 64×64 and use random rotation and horizontally flipping for data augmentation. The mini-batch size is set to 32 and total training iterations are set to 800,000. The learning rate is initialized as 2×10^{-4} . It remains constant for the first 270,000 iterations and then decreases to 10^{-6} in the next 560,000 iterations following the cosine annealing. They adopt the Adam optimizer ($\beta_1 = 0.9, \beta_2 = 0.9$) to train the model. During test, they first apply self-ensemble trick for each model, which could involves 8 outputs for fusion. Then, they fuse the self-ensembled outputs of the CIBDNet, CIBDNet_NF and CIBDNet_NFE models, respectively.

4.5 IVL Team




Fig. 5. The method proposed by the IVL Team.

The architecture proposed by the IVL team for the video challenge track is shown in Fig. 5 and contains three cascaded modules. The first module stacks the input frames (five consecutive frames are used) and extracts deep features from them. The second module aligns the features extracted from the adjacent frames with the features of the target frame. This is achieved by using a Spatio-Temporal Offset Prediction Network (STOPN), which implements a U-Net like architecture to estimate the deformable offsets that are later applied to deform a regular convolution and produce spatially-aligned features. Inspired by [20], STOPN predicts spatio-temporal offsets that are different at each spatial and temporal position. Moreover, as stated by [59], they apply deformable alignment at feature level to increase alignment accuracy. The third module contains two groups of standard and transposed convolutions to progressively perform feature fusion and upscaling. The input target frame is then $\times 4$ upscaled using bicubic interpolation and finally added to the network output to produce the final result. They only process the luma channel (Y) of the input frames because it contains the most relevant information on the scene details. The final result is obtained using the restored Y channel and the original chroma channels upsampled using bicubic interpolation, followed by a RGB conversion.

They train the model for 250,000 iterations using a batch size equal to 32. Patches with the size of 96×96 pixels are used, and data augmentation with random flip is applied. They set the temporal neighborhood N to five, hence they stack the target frame with the two previous and the two subsequent frames. The learning rate was initially set to 10^{-3} for the first 200,000 iterations, then reduced to 10^{-4} for the remaining iterations. They use MSE as the loss function and optimize it using the Adam optimizer.

4.6 SRC-B Team

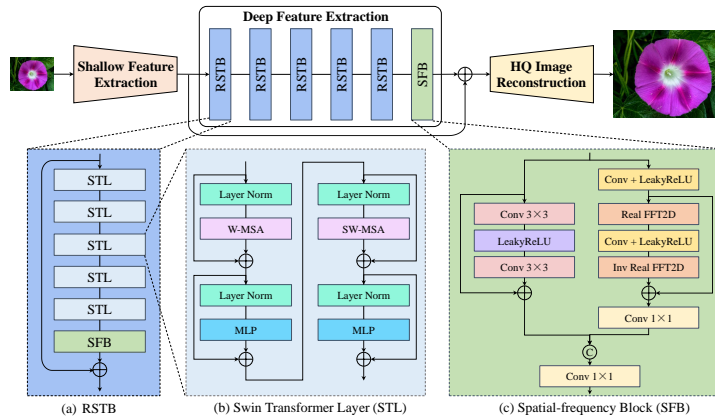


Fig. 6. The SwinFIR method proposed by the SRC-B Team.

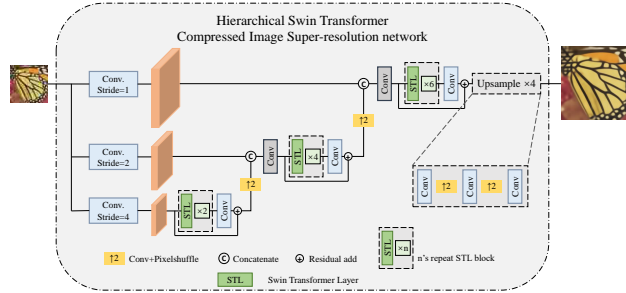


Fig. 7. Overview of HST method proposed by the USTC-IR Team. STL block is the Swin Transformer layer from SwinIR [48].

Inspired by SwinIR [48], the SRC-B Team proposes the SwinFIR method with the Swin Transformer [52] and the Fast Fourier Convolution [14]. As shown in Fig. 6, SwinFIR consists of three modules: shallow feature extraction, deep feature extraction and high-quality (HQ) image reconstruction modules. The shallow feature extraction and high-quality (HQ) image reconstruction modules adopt the same configuration as SwinIR. The residual Swin Transformerblock (RSTB) is a residual block with Swin Transformerlayers (STL) and convolutional layers in SwinIR. They all have local receptive fields and cannot extract the global information of the input image. The Fast Fourier Convolution has the ability to extract global features, so they replace the convolution (3×3) with Fast Fourier Convolution and a residual module to fuse global and local features, named Spatial-frequency Block (SFB), to improve the representation ability of model.

They use the Adam optimizer with default parameters and the Charbonnier L1 loss [43] to train the model. The initial learning rate is 2×10^{-4} , and they use the cosine annealing learning rate scheduler [53] with about 500,000 iterations. The batch size is 32 and patch size is 64. They use horizontal flip, vertical flip, rotation, RGB perm and mixup [79] for data augmentation.

4.7 USTC-IR Team [44]

The USTC-IR Team proposes a Hierarchical Swin Transformer (HST) for compressed image super-resolution, which is inspired by multi-scale-based frameworks [55,46,56] and transformer-based frameworks [48,47]. As shown in Fig. 7, the network is divided into three branches so that it can learn global and local information from different scales. Specifically, the input image is first downsampled to different scales by convolutions. Then, it is fed to different Residual Swin Transformer Blocks (RSTB) from SwinIR [48] to obtain the restored hierarchical features from each scale. To fuse the features from different scales, they super-

resolve the low-scale feature and concatenate it with the higher feature. Finally, there is a pixel-shuffle block to implement the $4\times$ super-resolution of features.

The training images are paired-cropped into 64×64 patches, and augmented by random horizontal flips, vertical flips and rotations. They train the model by the Adam optimizer with the initial learning rate of 2×10^{-4} . The learning rate is decayed by the factor of 0.5 twice, at the 200,000th and the 300,000 steps, respectively. The network is first trained by Charbonnier loss [43] for about 50,000 steps and finetuned by the MSE loss until convergence.

4.8 MSDRSR Team

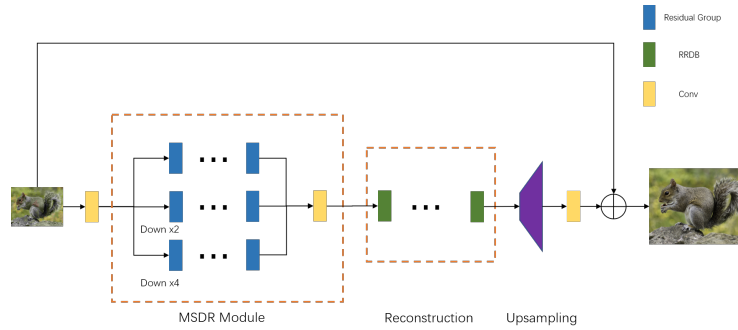


Fig. 8. The architecture of the method proposed by the MSDRSR Team. It utilizes a multi-scale degradation removal module, which employs the multi-scale structure to achieve balance between detail enhancement and artifacts removal.




Fig. 9. Enhanced Residual Group.

The architecture of the method proposed by the MSDRSR Team is illustrated in Fig. 8. The details are described in the following.

MSDR Module. In the MSDR Module, a Multi-Scale Degradation Removal (MSDR) module is employed after the first convolutional layer. The MSDR module uses several Enhanced Residual Group (ERG) to extract multi-scale features and can achieve a better trade-off between detail enhancement and compression

artifacts removal. The architecture of enhanced residual group (ERG) is illustrated in Fig. 9. ERG removes the channel attention module in each residual block, and adds a high-frequency attention block [24] at the end of the residual block. Compared with the original design of residual group, ERG can effectively remove artifacts while reconstructing high-frequency details. Moreover, our proposed ERG is very efficient and does not introduce much runtime overhead.

Reconstruction Module. The reconstruction module is build on ESRGAN [67], which has 23 residual-in-residual dense blocks (RRDB). To further improve the performance [49], they change activation function to SiLU [26].

The MSDRSR employs a two-stage training strategy. In each stage, MSDRSR is first trained with Laplacian pyramid loss with the patch size of 256, and then it is fine-tuned with the MSE loss with the patch size of 640. They augment the training data with random flipping and rotations. In the first stage, MSDRSR is trained on DF2K [1,63] for 100,000 iterations with the batch size of 64. It adopts the Adam optimizer with an initial learning rate of 5×10^{-4} . The Cosine scheduler is adopted with the minimal learning rate of 5×10^{-5} . Then, MSDRSR is fine-tuned with learning rate of 10^{-5} for 20,000 iterations. In the second stage, MSDRSR loads pre-trained weights from the first stage, and then they add 10 more randomly initialized blocks to the feature extractor. It is trained on DF2K and DIV8K [29] datasets. Then MSDRSR adopts the same training strategy as the first stage.

4.9 Giantpandacv

Inspired by previous image restoration and JPEG artifacts removal research [39,16], the Giantpandacv Team proposes the Artifact-aware Attention Network (A³Net) that can use the global semantic information of the image to adaptive control the trade-off between artifacts removal and details restored. Specifically, the A³Net uses an encoder to extract image texture features and artifact-aware features simultaneously, and then it adaptively removes image artifacts through a dynamic controller and a decoder. Finally, the A³Net uses some nonlinear-activation-free blocks to build a reconstructor to further recover the lost high-frequency information, resulting in a high resolution image.

The main architecture of the A³Net is shown in Figure. 10, which consists of four components: Encoder, Decoder, Dynamic Controller, and Reconstructor. The details of these modules of A³Net are described as follows.

Encoder: The Encoder aims to extract the deep features and decouple the latent artifact-aware features from the input image. The Encoder contains four scales, each of which has a skip connection to connect to the decoder. In order to improve the computational efficiency of the network, we use 2 Nonlinear Activation Free (NAF) blocks [16] at each scale. The number of output channels in each layer from the first to the fourth scale is set to 64, 128, 256, and 512, respectively. The image features from the encoder are passed into the decoder. At the same time, the global average pooling layer is used to get the artifact-aware features from the image features.

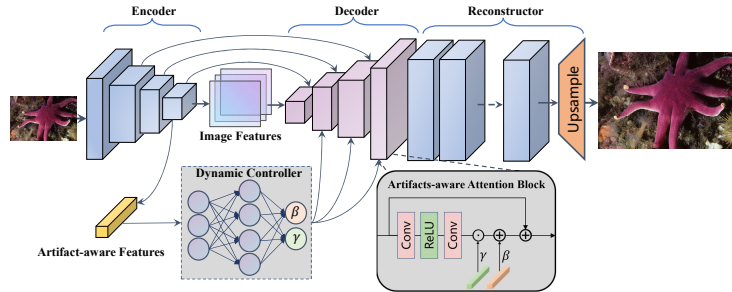


Fig. 10. Architecture of the A³Net proposed by the Giantpandacy Team.

Dynamic Controller: The dynamic controller is a 3-layer MLP and take as input the artifact-aware features, representing the latent degree of image compression. The main purpose of the dynamic controller is to allow the latent degree of image compression to be flexibly applied to the decoder, thus effectively removing artifacts. Inspired by recent research in spatial feature transform [57,66], we employ dynamic controller to generate modulation parameters pair (γ, β) which embed on the decoder. Moreover, we used three different final layers of the MLP to accommodate the different scales of features.

Decoder: The decoder consists of artifact-aware attention blocks with three different scales. The artifact-aware attention blocks mainly removes artifacts by combining image features and embedded artifact-aware parameters (γ, β) . The number of artifact-aware attention blocks in each scale is set to 4. It can be expressed as follows:

$$F_{out} = \gamma \odot F_{in} + \beta \quad (1)$$

where F_{in} and F_{out} denote the feature maps before and after the affine transformation, and \odot is referred to as element-wise multiplication.

Reconstructor: The aim of the reconstructor is to further restore the lost texture details, and then the features are up-sampled to reconstruct a high-resolution image. Specifically, they use a deeper NAF to facilitate the network capturing similar textures over long distances, thus obtaining more texture details.

Implementation and training details: The numbers of NAF blocks in the each scale of the Encoder and the Reconstructor are flexible and configurable, which are set to 2 and 8, respectively. For the up-scaling module, they use pixel-shuffle to reconstruct a high-resolution image. During training, the A³Net is trained on the crop training dataset with LR and HR pairs. The input pairs are randomly cropped to 512×512 . Random rotation and random horizontal flop are applied for data augmentation. They use the AdamW optimizer with the learning rate of 2×10^{-4} learning rate to train the model for 1,000,000 iterations and the learning rate is decayed with the cosine strategy. Weight decay is 10^{-4} for all the training periodic.

4.10 Aselsan Research Team




Fig. 11. The dual-domain super-resolution network of the Aselsan Research Team.

The Aselsan Research Team proposes a dual-domain super-resolution network, shown in Fig. 11. The network utilizes information in both pixel and wavelet domains. The information of these domains are processed in parallel by a modified IMDeception Network [2] to further increase the receptive field and the capability for processing non-local information. The two branches of the modified IMDeception network generates the super-resolved image and the enhanced low-resolution image, respectively. The super-resolved images are fused through a pixel attention network as used in [4] and enhanced low-resolution images are averaged for fusion. These low-resolution outputs are used during training to further guide the network and add the dual capability to the network. To further boost the performance of the entire network, the structure is encapsulated in a geometric ensembling architecture. We used LR_fused output as well as SR_fused output for training, using LR as guidance through out the optimization. The loss function we used is as follows

$$L = 0.5 \times L2(\text{LR_fused} - \text{Uncompressed_LR}) + 0.5 \times L2(\text{SR_fused} - \text{HR}) \quad (2)$$

Note that almost entire network is shared for this dual purpose, having a secondary and complementary guidance which is able to boost the performance by around 0.05dB.

They use the Adam optimizer with $\beta_1 = 0.9, \beta_2 = 0.999$ for training. The batch size is set to 8, and the training samples are cropped to 512×512 . The learning rate is initialized to 5^{-4} and it is decayed with the factor of 0.75 every 200 epochs (800 iteration in each epoch). The model is totally trained for 2,000 epochs.

4.11 SRMUI Team [17]

The method proposed by the SRMUI Team is illustrated in Fig. 12. They propose some modifications of SwinIR [48] (based on Swin Transformer [52]) that enhance

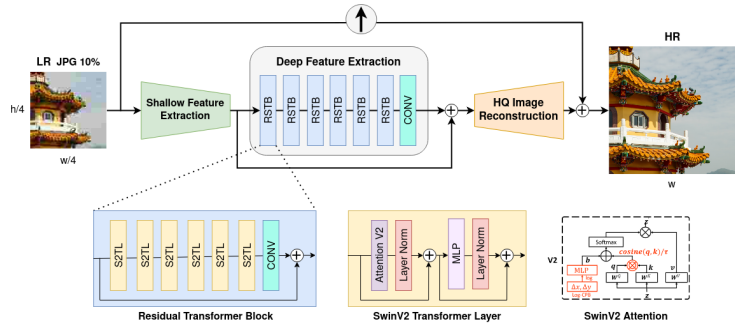


Fig. 12. The Swin2SR method of the SRMUI Team.

the model’s capabilities for super-resolution, and in particular, for compressed input SR. They update the original Residual Transformer Block (RSTB) by using the new SwinV2 transformer [51] layers and attention to scale up capacity and resolution. This method has a classical upscaling branch which uses a bicubic interpolation to recover basic structural information. The output of our model is added to the basic upscaled image to enhance it. They also explored different loss functions to make the model more robust to JPEG compression artifacts, being able to recover high-frequency details from the compressed LR image, and therefore it is able to achieve better performance.

4.12 MVideo Team

The MVideo Team proposes a two-stage network for compressed image super resolution. The overall architecture of the network is as Fig. 15. The Deblock-Net in first stage takes the compressed low resolution image with JPEG blocky artifacts as input and outputs the enhanced low resolution image. Then the SR-Net in the second stage is applied to the enhanced low resolution image, and generate the final SR image. Both networks use RRDBNet [67] as implementation, while they remove the pixel unshuffle operation from the first Deblock-Net

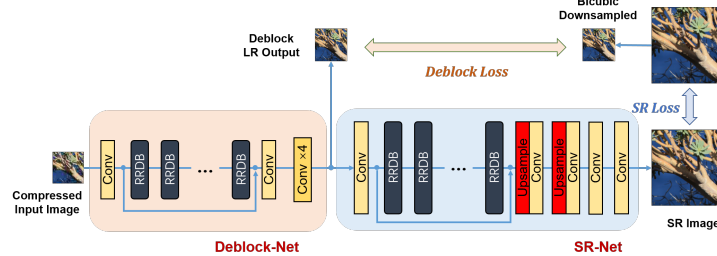


Fig. 13. The two-stage network proposed by the MVideo Team.

in the beginning and the upsample operation at last. The SR-Net uses the same hyper-parameters as used in ESRGAN [67].

Based on this pipeline, they train the two networks separately to reduce training time consumption. Firstly, they use the pretrained weights from the official ESRGAN, and load the SR-Net with it. Then the SR-Net is frozen and the deblock loss between bicubic downsampled LR from ground truth HR and deblock output is applied in order to train the Deblock-Net. After the training of the deblock net finishes, they use both deblock loss and SR loss (final output and ground truth) to train the model, with the weight of the deblock loss of 0.01 and the weight of the SR loss of 1.0. Then the model is finetuned only using the SR loss to improve the PSNR of the final output. Detailed training settings are listed below:

- (I) Pre-train the Deblock-Net. Firstly load the pretrained RRDBNet to SR-Net, and then use only deblock loss to train the Deblock-Net. The patch size is 128 and batch size is 32. Training is for 50k iterations using Adam optimizer, and learning rate is initiated with 5×10^{-4} , which decreases with the factor of 0.5 at the 30,000th and 40,000th iterations, respectively.
- (II) End-to-end training of the two-stage network. Using the pretrained weights from stage (I), they train the full network using both the deblock loss (weight=0.01) and the SR loss (weight=1.0), which mainly focuses on the learning of SR-Net. In this stage, patch size, batch size, and Adam optimizer stay the same as above, and they use the CosineAnnealingRestartLR scheduler with all periods of 50,000 for 200,000 iterations.
- (III) Finetuning the last weights using patch size of 512 and batch size of 8. The initial learning rate is set as 2×10^{-5} and the learning rate decays by the factor at the 20,000th, 30,000th, and 40,000th iterations, respectively. This stage takes totally 50,000 iterations.
- (IV) The last stage finetunes the model from the previous stage for 50,000 iterations, with the patch size of 256 and the batch size of 8, initial learning rate of 5×10^{-6} and multistep scheduler which decreases the learning rate by the factor of 0.5 at the 20,000th, 30,000th and 40,000th iterations, respectively. The final model is used in the inference phase.

4.13 UESTC+XJU CV Team




Fig. 14. Method of the UESTC+XJU CV Team: based on Restormer [80]. RDCB: A Channel-attention Block is inserted into the Residual Dense connection block.

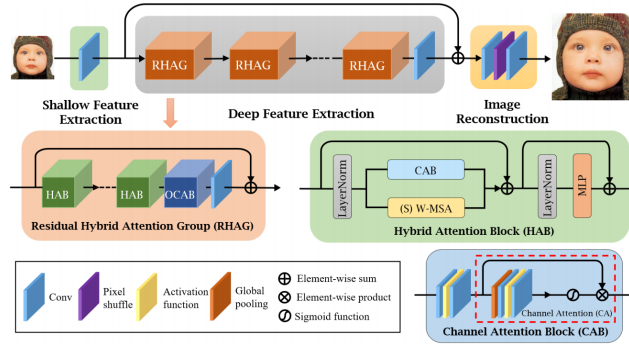


Fig. 15. The HAT [13] method used by the cvlab Team in Track 1.

The UESTC+XJU CV Team utilizes Restormer [80] for compressed image enhancement. The overall structure is shown in Fig. 14. First, the compressed image is input into the Restormer (the last layer of the original network is removed), and then upsampled by two convolutional layers and deconvolutional layers. Finally, the feature map is input into the common CNN network, which consists of L (4 in this model) Residual Dense Channel-attention Blocks (RD-CBs), and then it is added with the original upsampled image to obtain the reconstructed image. Specially, the channel-attention layer [83] is inserted after the four dense connection layers in the residual dense block [84].

In the training process, the raw image is cropped into patches with the size of 256×256 as the training samples, and the batch size is set to 8. They also adopt flip and rotation as data augmentation strategies to further expand the dataset. The model is trained by Adam optimizer [41] and cosine annealing learning rate scheduler for 3×10^5 iterations. The learning rate is initially set to 3×10^{-4} . They use L2 loss as the loss function.

4.14 cvlab Team

The cvlab Team uses HAT [13] as the solution in Track 1. As show in Fig. 15, The number of RHAG blocks is set to 6. The number of HAB blocks in each RHAG block is 6. The number of feature channels is 180. In the training process, the raw and compressed sequences are cropped into 64×64 patches as the training pairs, and the batch size is set to 4. They also adopt flip and rotation as data augmentation strategies to further expand the dataset. We train all models using Adam [41] optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 10^{-6}$, and the learning rate is initially set to 2×10^{-4} and decays linearly to 5×10^{-5} after 200,000 iterations, which keeps unchanged until 400,000 iterations. Then, the learning rate is further decayed to 2×10^{-5} and 1×10^{-5} until converged. The total number of iterations is 1,000,000. They use L1 loss as the loss function.

Acknowledgments

We thank the sponsors of the AIM and Mobile AI 2022 workshops and challenges: AI Witchlabs, MediaTek, Huawei, Reality Labs, OPPO, Synaptics, Raspberry Pi, ETH Zürich (Computer Vision Lab) and University of Würzburg (Computer Vision Lab).

Appendix: Teams and affiliations

AIM 2022 Team

Challenge:

AIM 2022 Challenge on Super-Resolution of Compressed Image and Video

Organizer(s):

Ren Yang¹ (ren.yang@vision.ee.ethz.ch),

Radu Timofte^{1,2} (radu.timofte@uni-wuerzburg.ch)

Affiliation(s):

¹ Computer Vision Lab, ETH Zürich, Switzerland

² Julius Maximilian University of Würzburg, Germany

VUE Team

Member(s):

Xin Li¹ (lixin41@baidu.com), Qi Zhang¹, Lin Zhang², Fanglong Liu¹, Dongliang He¹, Fu li¹, He Zheng¹, Weihang Yuan¹

Affiliation(s):

¹ Department of Computer Vision Technology (VIS), Baidu Inc.

² Institute of Automation, Chinese Academy of Sciences

NoahTerminalCV Team

Member(s):

Pavel Ostyakov (ostyakov.pavel@huawei.com), Dmitry Vyal, Magauiya Zhussip, Xueyi Zou, Youliang Yan

Affiliation(s):

Noah's Ark Lab, Huawei

BSR Team

Member(s):

Lei Li (lilei.leili@bytedance.com), Jingzhu Tang, Ming Chen, Shijie Zhao

Affiliation(s):

Multimedia Lab, ByteDance Inc.

CASIA LCVG Team

Member(s):

Yu Zhu¹ (zhuyu.cv@gmail.com), Xiaoran Qin¹, Chenghua Li¹, Cong Leng^{1,2,3}, Jian Cheng^{1,2,3}

Affiliation(s):

¹ Institute of Automation, Chinese Academy of Sciences, Beijing, China

² MAICRO, Nanjing, China

³ AiRiA, Nanjing, China

IVL Team

Member(s):

Claudio Rota (c.rota30@campus.unimib.it), Marco Buzzelli, Simone Bianco, Raimondo Schettini

Affiliation(s):

University of Milano - Bicocca, Italy

Samsung Research China – Beijing (SRC-B)

Member(s):

Dafeng Zhang (dfeng.zhang@samsung.com), Feiyu Huang, Shizhuo Liu, Xiaobing Wang, Zhezhu Jin

Affiliation(s):

Samsung Research China – Beijing (SRC-B), China

USTC-IR

Member(s):

Bingchen Li (1bc31415926@mail.ustc.edu.cn), Xin Li

Affiliation(s):

University of Science and Technology of China, Hefei, China

MSDRSR

Member(s):

Mingxi Li (li_mx_0318@163.com), Ding Liu¹

Affiliation(s):

¹ ByteDance Inc.

Giantpandacv Team

Member(s):

Wenbin Zou^{1,4} (alexzou14@foxmail.com), Peijie Dong², Tian Ye³, Yunchen Zhang⁵, Ming Tan⁴, Xin Niu²

Affiliation(s):

¹ South China University of Technology, Guangzhou, China

² National University of Defense Technology, Changsha, China

³ Jimei University, Xiamen, China

⁴ Fujian Normal University, Fuzhou, China

⁵ China Design Group Inc., Nanjing, China

Aselsan Research Team

Member(s):

Mustafa Ayazoglu (mayazoglu@aselsan.com.tr)

Affiliation(s):

Aselsan (www.aselsan.com.tr), Ankara, Turkey

SRMUI Team

Member(s):

Marcos V. Conde¹ (marcos.conde-osorio@uni-wuerzburg.de), Ui-Jin Choi², Radu Timofte¹

Affiliation(s):

¹ Computer Vision Lab, Julius Maximilian University of Würzburg, Germany ² MegaStudyEdu, South Korea

MVideo Team

Member(s):

Zhuang Jia (jiazhuang@xiaomi.com), Tianyu Xu, Yijian Zhang

Affiliation(s):

Xiaomi Inc.

UESTC+XJU CV Team

Member(s):

Mao Ye (cvlab.uestc@gmail.com), Dengyan Luo, Xiaofeng Pan

Affiliation(s):

University of Electronic Science and Technology of China, Chengdu, China

cvlab Team**Member(s):**Liuhan Peng¹ (pengliuhan@gmail.com), Mao Ye²**Affiliation(s):**¹ Xinjiang University, Xinjiang, China² University of Electronic Science and Technology of China, Chengdu, China**References**

1. Agustsson, E., Timofte, R.: NTIRE 2017 challenge on single image super-resolution: Dataset and study. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 126–135 (2017) [2](#), [14](#)
2. Ayazoğlu, M.: Imdeception: Grouped information distilling super-resolution network. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 756–765 (2022) [16](#)
3. Bhat, G., Danelljan, M., Timofte, R.: NTIRE 2021 challenge on burst super-resolution: Methods and results. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 613–626 (2021) [8](#)
4. Bilecen, B.B., Fişne, A., Ayazoğlu, M.: Efficient multi-purpose cross-attention based image alignment block for edge devices. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 3639–3648 (2022) [16](#)
5. Boytsov, L., Naidan, B.: Engineering efficient and effective non-metric space library. In: Brisaboa, N.R., Pedreira, O., Zezula, P. (eds.) Similarity Search and Applications - 6th International Conference, SISAP 2013, A Coruña, Spain, October 2-4, 2013, Proceedings. Lecture Notes in Computer Science, vol. 8199, pp. 280–293. Springer (2013). https://doi.org/10.1007/978-3-642-41062-8_28, https://doi.org/10.1007/978-3-642-41062-8_28 [8](#)
6. Briechele, K., Hanebeck, U.D.: Template matching using fast normalized cross correlation. In: Optical Pattern Recognition XII. vol. 4387, pp. 95–102. International Society for Optics and Photonics (2001) [8](#)
7. Bychkovsky, V., Paris, S., Chan, E., Durand, F.: Learning photographic global tonal adjustment with a database of input/output image pairs. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2011) [5](#)
8. Caballero, J., Ledig, C., Aitken, A., Acosta, A., Totz, J., Wang, Z., Shi, W.: Real-time video super-resolution with spatio-temporal networks and motion compensation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4778–4787 (2017) [2](#)
9. Chan, K.C., Wang, X., Yu, K., Dong, C., Loy, C.C.: 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 (CVPR) (2021) [2](#)
10. Chan, K.C., Zhou, S., Xu, X., Loy, C.C.: Basicvsr++: Improving video super-resolution with enhanced propagation and alignment. arXiv preprint arXiv:2104.13371 (2021) [2](#), [7](#)

11. Chen, P., Yang, W., Wang, M., Sun, L., Hu, K., Wang, S.: Compressed domain deep video super-resolution. *IEEE Transactions on Image Processing* **30**, 7156–7169 (2021) [2](#)
12. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: *International Conference on Machine Learning (ICML)*. pp. 1597–1607. PMLR (2020) [8](#)
13. Chen, X., Wang, X., Zhou, J., Dong, C.: Activating more pixels in image super-resolution transformer. *arXiv preprint arXiv:2205.04437* (2022) [9](#), [10](#), [19](#)
14. Chi, L., Jiang, B., Mu, Y.: Fast fourier convolution. *Advances in Neural Information Processing Systems (NeurIPS)* **33**, 4479–4488 (2020) [12](#)
15. Chu, X., Chen, L., , Chen, C., Lu, X.: Improving image restoration by revisiting global information aggregation. *arXiv preprint arXiv:2112.04491* (2021) [5](#)
16. Chu, X., Chen, L., Yu, W.: NAFSSR: Stereo image super-resolution using nafnet. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1239–1248 (2022) [14](#)
17. Conde, M.V., Choi, U.J., Burchi, M., Timofte, R.: Swin2SR: Swinv2 transformer for compressed image super-resolution and restoration. In: *Proceedings of the European Conference on Computer Vision (ECCV) Workshops* (2022) [5](#), [16](#)
18. Conde, M.V., Timofte, R., et al.: Reversed image signal processing and raw reconstruction. AIM 2022 challenge report. In: *Proceedings of the European Conference on Computer Vision (ECCV) Workshops* (2022) [2](#)
19. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: A large-scale hierarchical image database. In: *2009 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 248–255 (2009) [5](#), [10](#)
20. Deng, J., Wang, L., Pu, S., Zhuo, C.: Spatio-temporal deformable convolution for compressed video quality enhancement. *Proceedings of the AAAI Conference on Artificial Intelligence* **34**(07), 10696–10703 (2020) [11](#)
21. Deng, X., Yang, R., Xu, M., Dragotti, P.L.: Wavelet domain style transfer for an effective perception-distortion tradeoff in single image super-resolution. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 3076–3085 (2019) [2](#)
22. Dong, C., Deng, Y., Loy, C.C., Tang, X.: Compression artifacts reduction by a deep convolutional network. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. pp. 576–584 (2015) [2](#)
23. Dong, C., Loy, C.C., He, K., Tang, X.: Learning a deep convolutional network for image super-resolution. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 184–199. Springer (2014) [2](#)
24. Du, Z., Liu, D., Liu, J., Tang, J., Wu, G., Fu, L.: Fast and memory-efficient network towards efficient image super-resolution. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 853–862 (2022) [14](#)
25. Ehrlich, M., Davis, L., Lim, S.N., Shrivastava, A.: Quantization guided jpeg artifact correction. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 293–309. Springer (2020) [2](#)
26. Elfving, S., Uchibe, E., Doya, K.: Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural Networks* **107**, 3–11 (2018) [14](#)
27. Google: YouTube. <https://www.youtube.com> [3](#), [4](#), [6](#)
28. Gu, J., Lu, H., Zuo, W., Dong, C.: Blind super-resolution with iterative kernel correction. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1604–1613 (2019) [2](#)

29. Gu, S., Lugmayr, A., Danelljan, M., Fritsche, M., Lamm, J., Timofte, R.: DIV8K: Diverse 8k resolution image dataset. In: 2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW). pp. 3512–3516 (2019) [5](#), [14](#)
30. Guan, Z., Xing, Q., Xu, M., Yang, R., Liu, T., Wang, Z.: MFQE 2.0: A new approach for multi-frame quality enhancement on compressed video. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2019) [2](#)
31. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 770–778 (2016) [8](#)
32. Ignatov, A., Timofte, R., Denna, M., Younes, A., et al.: Efficient and accurate quantized image super-resolution on mobile npus, mobile AI & AIM 2022 challenge: Report. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops (2022) [2](#)
33. Ignatov, A., Timofte, R., Kuo, H.K., Lee, M., Xu, Y.S., et al.: Real-time video super-resolution on mobile npus with deep learning, mobile AI & AIM 2022 challenge: Report. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops (2022) [2](#)
34. Ignatov, A., Timofte, R., et al.: Efficient bokeh effect rendering on mobile gpus with deep learning, mobile AI & AIM 2022 challenge: Report. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops (2022) [2](#)
35. Ignatov, A., Timofte, R., et al.: Efficient single-image depth estimation on mobile devices, mobile AI & AIM challenge: Report. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops (2022) [2](#)
36. Ignatov, A., Timofte, R., et al.: Learned smartphone isp on mobile gpus with deep learning, mobile AI & AIM 2022 challenge: Report. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops (2022) [2](#)
37. Isobe, T., Jia, X., Gu, S., Li, S., Wang, S., Tian, Q.: Video super-resolution with recurrent structure-detail network. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 645–660. Springer (2020) [2](#)
38. Jiang, J., Zhang, K., Timofte, R.: Towards flexible blind jpeg artifacts removal. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 4997–5006 (2021) [2](#)
39. Jiang, J., Zhang, K., Timofte, R.: Towards flexible blind jpeg artifacts removal. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 4997–5006 (2021) [14](#)
40. Kim, J., Lee, J.K., Lee, K.M.: Accurate image super-resolution using very deep convolutional networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1646–1654 (2016) [2](#)
41. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Proceedings of the International Conference on Learning Representations (ICLR) (2015) [19](#)
42. Kınık, F.O., Menteş, S., Özcan, B., Kirac, F., Timofte, R., et al.: AIM 2022 challenge on instagram filter removal: Methods and results. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops (2022) [2](#)
43. Lai, W.S., Huang, J.B., Ahuja, N., Yang, M.H.: Fast and accurate image super-resolution with deep laplacian pyramid networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **41**(11), 2599–2613 (2018) [12](#), [13](#)
44. Li, B., Li, X., Lu, Y., Liu, S., Feng, R., Chen, Z.: HST: Hierarchical swin transformer for compressed image super-resolution. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops (2022) [5](#), [12](#)

45. Li, L., Tang, J., Chen, M., Zhao, S., Li, J., Zhang, L.: Multi-patch learning: Looking more pixels in the training phase. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops (2022) [5](#), [8](#)
46. Li, X., Sun, S., Zhang, Z., Chen, Z.: Multi-scale grouped dense network for vvc intra coding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 158–159 (2020) [12](#)
47. Liang, J., Cao, J., Fan, Y., Zhang, K., Ranjan, R., Li, Y., Timofte, R., Van Gool, L.: VRT: A video restoration transformer. arXiv preprint arXiv:2201.12288 (2022) [2](#), [12](#)
48. Liang, J., Cao, J., Sun, G., Zhang, K., Van Gool, L., Timofte, R.: Swinir: Image restoration using swin transformer. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 1833–1844 (2021) [2](#), [6](#), [12](#), [16](#)
49. Lin, Z., Garg, P., Banerjee, A., Magid, S.A., Sun, D., Zhang, Y., Van Gool, L., Wei, D., Pfister, H.: Revisiting rcan: Improved training for image super-resolution. arXiv preprint arXiv:2201.11279 (2022) [14](#)
50. Liu, H., Ruan, Z., Zhao, P., Dong, C., Shang, F., Liu, Y., Yang, L., Timofte, R.: Video super-resolution based on deep learning: a comprehensive survey. Artificial Intelligence Review pp. 1–55 (2022) [2](#)
51. Liu, Z., Hu, H., Lin, Y., Yao, Z., Xie, Z., Wei, Y., Ning, J., Cao, Y., Zhang, Z., Dong, L., et al.: Swin transformer v2: Scaling up capacity and resolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 12009–12019 (2022) [17](#)
52. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin Transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 10012–10022 (2021) [12](#), [16](#)
53. Loshchilov, I., Hutter, F.: SGDR: Stochastic gradient descent with warm restarts. arXiv preprint arXiv:1608.03983 (2016) [12](#)
54. Malkov, Y.A., Yashunin, D.A.: Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. CoRR [abs/1603.09320](#) (2016), <http://arxiv.org/abs/1603.09320> [8](#)
55. Pang, Y., Li, X., Jin, X., Wu, Y., Liu, J., Liu, S., Chen, Z.: FAN: Frequency aggregation network for real image super-resolution. In: European Conference on Computer Vision (ECCV). pp. 468–483. Springer (2020) [12](#)
56. Papyan, V., Elad, M.: Multi-scale patch-based image restoration. IEEE Transactions on Image Processing **25**(1), 249–261 (2015) [12](#)
57. Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2337–2346 (2019) [15](#)
58. Qin, X., Zhu, Y., Li, C., Wang, P., Cheng, J.: CIDBNet: A consecutively-interactive dual-branch network for jpeg compressed image super-resolution. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops (2022) [5](#), [9](#)
59. Rota, C., Buzzelli, M., Bianco, S., Schettini, R.: Video restoration based on deep learning: a comprehensive survey (2022) [11](#)
60. Sajjadi, M.S., Scholkopf, B., Hirsch, M.: Enhancenet: Single image super-resolution through automated texture synthesis. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 4491–4500 (2017) [2](#)
61. Tai, Y., Yang, J., Liu, X., Xu, C.: Memnet: A persistent memory network for image restoration. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 4539–4547 (2017) [2](#)

62. Tao, X., Gao, H., Liao, R., Wang, J., Jia, J.: Detail-revealing deep video super-resolution. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 4472–4480 (2017) [2](#)
63. Timofte, R., Agustsson, E., Van Gool, L., Yang, M.H., Zhang, L.: NTIRE 2017 challenge on single image super-resolution: Methods and results. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 114–125 (2017) [4](#), [5](#), [14](#)
64. Timofte, R., Rothe, R., Van Gool, L.: Seven ways to improve example-based single image super resolution. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1865–1873 (2016) [4](#)
65. Wang, X., Chan, K.C., Yu, K., Dong, C., Change Loy, C.: EDVR: Video restoration with enhanced deformable convolutional networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops (2019) [2](#)
66. Wang, X., Yu, K., Dong, C., Loy, C.C.: Recovering realistic texture in image super-resolution by deep spatial feature transform. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 606–615 (2018) [15](#)
67. Wang, X., Yu, K., Wu, S., Gu, J., Liu, Y., Dong, C., Qiao, Y., Change Loy, C.: ESR-GAN: Enhanced super-resolution generative adversarial networks. In: Proceedings of the European Conference on Computer Vision Workshops (ECCVW). pp. 0–0 (2018) [14](#), [17](#), [18](#)
68. Xu, Y., Gao, L., Tian, K., Zhou, S., Sun, H.: Non-local ConvLSTM for video compression artifact reduction. In: Proceedings of The IEEE International Conference on Computer Vision (ICCV) (October 2019) [2](#)
69. Yamac, M., Ataman, B., Nawaz, A.: Kernelnet: A blind super-resolution kernel estimation network. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 453–462 (2021) [2](#)
70. Yang, R., Mentzer, F., Gool, L.V., Timofte, R.: Learning for video compression with hierarchical quality and recurrent enhancement. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6628–6637 (2020) [2](#)
71. Yang, R., Sun, X., Xu, M., Zeng, W.: Quality-gated convolutional LSTM for enhancing compressed video. In: Proceedings of the IEEE International Conference on Multimedia and Expo (ICME). pp. 532–537. IEEE (2019) [2](#)
72. Yang, R., Timofte, R., et al.: NTIRE 2021 challenge on quality enhancement of compressed video: Dataset and study. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops (2021) [2](#)
73. Yang, R., Timofte, R., et al.: NTIRE 2021 challenge on quality enhancement of compressed video: Methods and results. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops (2021) [2](#)
74. Yang, R., Timofte, R., et al.: AIM 2022 challenge on super-resolution of compressed image and video: Dataset, methods and results. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops (2022) [2](#)
75. Yang, R., Timofte, R., et al.: NTIRE 2022 challenge on super-resolution and quality enhancement of compressed video: Dataset, methods and results. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops (2022) [2](#), [3](#)
76. Yang, R., Xu, M., Liu, T., Wang, Z., Guan, Z.: Enhancing quality for HEVC compressed videos. IEEE Transactions on Circuits and Systems for Video Technology (2018) [2](#)

77. Yang, R., Xu, M., Wang, Z.: Decoder-side HEVC quality enhancement with scalable convolutional neural network. In: Proceedings of the IEEE International Conference on Multimedia and Expo (ICME). pp. 817–822. IEEE (2017) [2](#)
78. Yang, R., Xu, M., Wang, Z., Li, T.: Multi-frame quality enhancement for compressed video. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6664–6673 (2018) [2](#)
79. Yoo, J., Ahn, N., Sohn, K.A.: Rethinking data augmentation for image super-resolution: A comprehensive analysis and a new strategy. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8375–8384 (2020) [12](#)
80. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H.: Restormer: Efficient transformer for high-resolution image restoration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5728–5739 (2022) [18](#), [19](#)
81. Zhang, K., Liang, J., Van Gool, L., Timofte, R.: Designing a practical degradation model for deep blind image super-resolution. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 4791–4800 (2021) [2](#)
82. Zhang, K., Zuo, W., Chen, Y., Meng, D., Zhang, L.: Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE Transactions on Image Processing* **26**(7), 3142–3155 (2017) [2](#)
83. Zhang, Y., Li, K., Li, K., Wang, L., Zhong, B., Fu, Y.: Image super-resolution using very deep residual channel attention networks. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 286–301 (2018) [2](#), [19](#)
84. Zhang, Y., Tian, Y., Kong, Y., Zhong, B., Fu, Y.: Residual dense network for image super-resolution. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2472–2481 (2018) [2](#), [6](#), [19](#)
85. Zheng, M., Xing, Q., Qiao, M., Xu, M., Jiang, L., Liu, H., Chen, Y.: Progressive training of a two-stage framework for video restoration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 1024–1031 (2022) [2](#), [4](#), [5](#), [6](#)