



OPTIMIZING UHD PRODUCTION WITH SUPER RESOLUTION TECHNIQUES

X. Ducloux¹, Jon Frydensbjerg² and M. Kolle²

¹ Harmonic, USA and ² Pixop, Denmark

ABSTRACT

Super resolution is a technique that is getting increased attention. It was first deployed in high-end TVs to upconvert content, and now content providers are looking at how it can be used on the production side to guarantee the same high quality on all devices. Different techniques can be used to apply super resolution to SD and HD content for transmitting UHD resolution. During a recent trial with France TV at the French Open, SD content with a 16:9 aspect ratio was processed offline to create UHD content, using a hybrid pipeline of deep restoration and super resolution methods. A pre-trained deep neural network was used to produce an intermediate upscaled and deblurred HD version, which was then passed on to a faster, less memory-intensive, machine-learning-based super resolution filter to produce the final UHD 4K resolution output. The result is a significantly better subjective viewing experience in terms of details and sharpness compared with classical filters used in live systems.

INTRODUCTION

Super resolution (SR) is a generic name to describe techniques applied on still images or videos to enhance the resolution and achieve the same quality as what can be achieved with a device capable of capturing higher resolution. Indeed, these techniques are much more than interpolation filters used in classical upscaling and can guess the missing information in the input source. SR has been successfully implemented in general image processing and microscopy over the last decade, thanks to the progress made in artificial intelligence. Many websites throughout the world address the market of photo restoration and enhancement. SR techniques are now emerging for the consumer and professional video market. SR was first deployed in high-end TVs to replace classical interpolation in the upconversion process. Content providers are now looking at how it can be used on the production side to guarantee the same high quality on all devices. Another need for SR arose during the COVID-19 global health crisis: with the lack of new live content and the increase in video consumption due to home confinement, broadcasters and telco operators had to dig out old footage from their catalogs. In doing so, they discovered that consumers have an appetite for classics (i.e., legendary sport events, concerts, films and TV series). This paper will explain how the SR associated with deep restoration techniques gives a second life to this type of content on modern and big TV screens.

The first section of the paper will review various machine learning (ML) and deep learning (DL) techniques that can be applied to produce 4K SR from SD and HD video content, with the pros, cons and best application usages.



The second section will describe an experiment done with France TV at the French Open tennis tournament in 2020 with a hybrid pipeline of deep restoration of SD content and upscale to 4K with SR. It will describe the performance of the solution tested in terms of visual quality as well as processing power requirement.

The last section will present how SR and deep restoration techniques can support live distribution systems and why now is the perfect time to start migrating the upconversion of video content from end-users' TVs to the headend.

MACHINE LEARNING AND DEEP LEARNING TECHNIQUES FOR SUPER RESOLUTION

SR is a popular, yet challenging computer vision method due to its underdetermined nature and because hundreds of techniques have been proposed by the scientific community over the years. However, most of these methods are impractical for real-world implementation when state-of-the-art accuracy is desired due to their runtime speed characteristics. Only a small number of techniques offer a feasible trade-off between engineering complexity, accuracy and runtime speed for processing video, which are covered in this section. For a full survey of the wide range approaches on single images before the deep learning era began, refer to Nasrollahi and Moeslund (1) and also Liu et al. (2) for a review of recent video super resolution (VSR) DL methods.

To achieve satisfactory performance with SR, real-world degradations must be adequately modelled. While some of these degradations lead to irreversible information loss, others can be addressed algorithmically. In summary, the main sources of degradation involved in video acquisition are classically modelled as: motion warping, blurring, down-sampling and noise. Additional degradation is introduced if progressive video is converted to interlaced form, and at video encoding time when a lossy compression technique, such as the popular quantized DCT blocking scheme, is used.

In general, single image SR (SISR) methods process a single image at a time, while video VSR algorithms deal with multiple successive images and frames at a time so as to utilize the relationship within frames to super-resolve the target frame. In a broad sense, VSR can be regarded as a class of SISR and is able to be processed by image SR algorithms frame by frame. However, the SR performance is not always satisfactory as artifacts may be brought in, which causes unguaranteed coherence within frames.

Learning-Based SR Algorithms

Learning-based or hallucination algorithms were first introduced in 1985 and use a neural network to improve the resolution of fingerprint images. These algorithms contain a training step in which the relationship between some high-resolution (HR) examples and their low-resolution (LR) counterparts are learned. This learned knowledge is then incorporated into the a priori term of the reconstruction. Importantly, the training database of learning-based SR algorithms needs to have a proper generalization capability.

Pixop Super Resolution

Pixop Super Resolution (PSR) is a VSR method inspired by RAISR, the work of Romano et al. (3). RAISR is an image processing technique that incorporates ML in order to produce



high-quality versions of low-resolution videos to save transmission bandwidth on Google+ for mobile devices. Its speed, image quality, and versatility design trade-off makes RAISR a really interesting candidate for video SR, as well.

PSR filters are trained on patches from corresponding low-res and high-res image pairs, according to three edge features:

- **Angle**, the direction of an edge
- **Strength**, the edge sharpness
- **Coherence**, a measure of edge directionality

A lookup hash table is then generated based on these features, the modulo of the pixel position and a combination of least-squares solvers.

Generating filters is essentially based on learning the least-squares solution that maps pixel patches from a LR input space to a HR output space. Image gradients are used to improve the image quality by observing that the best solution to processing edges is different from softer regions.

A visualization of PSR 3x filter kernels for a specific pixel position, learned from about 44 million patches of 11x11 pixels can be found in Figure 1. The multithreaded training process, which includes some GPU offloading, takes about 10 minutes. Notice how, from left to right, that filters have a selective response to the underlying edge being reconstructed. As the angle of the edge changes, the gradient angle of the filter rotates accordingly. Similarly, as the strength increases, the sharpness of the filters increases, and the directionality of the filter increases with rising coherence.

At runtime PSR selects and applies the most relevant learned filter via 2D convolution. When these filters are applied, they recreate details that are comparable to the original (while also avoiding aliasing issues like Moire patterns and jaggies), which is a significant improvement to interpolation. For more details about the technology, please refer to (4).

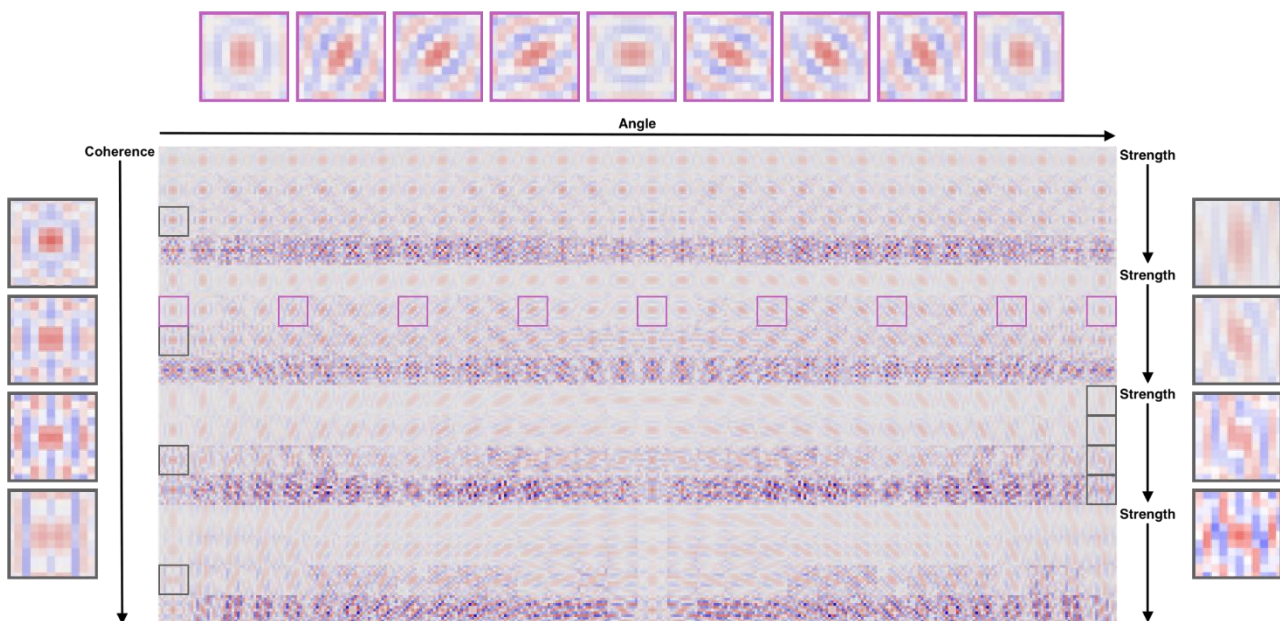


Figure 1 - A small subset of learned 2D convolution filter kernels for 3x super resolution. There are 32 angular steps, four edge strengths and four coherence levels, which makes for a total of 512 11x11 pixels kernels visualized. The intensity and color of



each kernel value in this visualization designates the weight and sign (red for positive feedback, blue for negative feedback).

Deep Learning Super Resolution Algorithms

Deep learning is a class of ML algorithms that uses two or more hidden neural network layers to progressively extract higher-level features from the raw input. For example, in image processing, lower layers may identify edges, while higher layers may identify the concepts relevant to a human such as digits or letters or faces. However, key foundational constructs related to DL such as feed-forward neural networks and the backpropagation algorithm used during training of deep neural networks will not be described here.

Since VSR techniques are an ongoing research topic in the scientific community, at this point they are far from a “solved” problem. Although great progress has been made by state-of-the-art VSR methods based on DL, especially on some public benchmark datasets, there are still many unresolved challenges when it comes to applying these methods in a real-world setting according to Liu et al. (2). For example, better evaluation criteria for video quality are needed to optimize the output so that it corresponds to how humans actually perceive video quality (as a replacement for PSNR and SSIM).

Autoencoders

An autoencoder is a neural network that learns to copy its input to its output. It has an internal (i.e., hidden) layer that describes a code used to represent the input, and it is constituted by two main parts: an encoder that maps the input into the code, and a decoder that maps the code to a reconstruction of the input. Autoencoders are usually restricted in ways that force them to compress the input representation, preserving only the most relevant aspects of the data in the copy. The idea of autoencoders has been popular in the field of neural networks for decades. The first applications date to the 1980s. Their most traditional application was dimensionality reduction or feature learning, but the autoencoder concept became more widely used for learning generative models of data.

A particularly interesting instance of a convolutional autoencoder-like architecture is the “UNet” which has become popular in the deep learning era. This type of neural network is also called “an hourglass architecture.” While originally used for image segmentation, UNet and other neural networks architectures that use the same approach show highly accurate results in a wide area of applications, including SR.

Generative Adversarial Networks

Generative Adversarial Networks (GANs) are an ML technique that offers excellent performance for a variety of tasks related to image generation. GANs solve a key challenge in generative models by introducing an additional neural network to score the quality of a generated output. This scorer neural network (called the discriminator) will score how realistic the image output by the generator neural network is. These two neural networks have opposing objectives (hence, the word adversarial). The generator network’s objective is to generate fake images that look real, while the discriminator network’s objective is to tell apart fake images from real ones.

This puts generative tasks in a setting where we have an ML model improving continuously by engaging in what is essentially a two-player game. As opposed to games like chess, the roles of the two players are asymmetric. In a GAN setting, one generates fake samples, and



the other distinguishes real ones from fake ones. Training drives the discriminator to attempt to learn to correctly classify samples as real or fake. Simultaneously, the generator attempts to fool the classifier into believing its samples are real. At convergence, the generator's samples are indistinguishable from real data, and the discriminator is discarded at inference time.

A key challenge with GANs is that they are time consuming to train. Non-convergence and “mode collapse” are issues to look out for, where the GAN fails to generalize properly, missing entire modes from the input data. For example, a GAN trained on the MNIST dataset (containing many samples of digits between 0 and 9), might completely omit a subset of the digits from its output.

Pixop Deep Restoration

Pixop Deep Restoration (PDR) is a VSR method designed to super-resolve and restore an HD version from a degraded SD input, such as an XDCAM encoding of digitally produced content.

Inspired by the work of Thomas (5), a single image, UNet SR algorithm was trained to hallucinate an image of twice the original resolution and mitigate the effects of degradations from gaussian noise, compression and blurring. The advantage of using a UNet is that it is much easier to train than a GAN, and by using various tricks such as introducing a perceptual loss function that utilizes a pretrained VGG-19 network, it is possible to achieve a visual output quality that resembles a GAN.

The model is based on a deep ResNet-34 UNet architecture that is trained on crop outs from a mix of about 10,000 frames from publicly available video and still image datasets, with the intention that it learns how to distinguish motion blur from the blur that results from upscaling an image. During the learning phase, the CNN is presented with thousands of image pairs of artificially degraded and perfect image patches. These degradations have been carefully engineered to resemble the type of artifacts commonly found in lossy compressed digital SD video. The training takes about three hours when implemented in the fastai v1 framework using the CUDA-accelerated version of PyTorch 1.0.1 running on a Nvidia GTX 1080 GPU.

At runtime video is processed frame by frame using only a single video frame as the input to generate the enhanced output in two steps:

1. The input frame is scaled to the desired resolution, which is typically anywhere from SD to HD.
2. An enhanced frame is produced from the output of the previous step via inference using our pretrained neural network model.

For more details about the technology, please refer to (6).

Hybrid Schemes Combining Deep Learning and Machine Learning

To upscale content from SD to 4K, we propose a pragmatic approach to VSR that takes advantage of the best features of the PSR and PDR algorithms.

Both methods are essentially single-image SR methods applied to video. However, while PSR is an accurate, memory-efficient and relatively performant SR upscaler even on a multi-core CPU, it is not designed for restoring degraded footage and therefore cannot achieve good results with older, low-quality SD footage by itself. This is not a training issue but an



inherent limitation due to the restricted number of parameters learned. PDR addresses these deficiencies with a higher algorithmic complexity, but it is limited to 1080 lines HD output due to GPU memory usage.

We therefore suggest the following hybrid pipeline:

1. Super-resolve and restore SD to 1080 lines HD via PDR
2. Upscale the restored HD from the first step to produce the final UHD 4K output via PSR

For the best results, the source is assumed to be clean by the restoration algorithm; therefore, noise reduction must typically be applied as a pre-processing step. Failure to provide a clean source results in visual artifacts because the pretrained model does not account for additive noise. For the same reason it is important that scanlines are properly aligned correctly within a frame, and if a digitized version of an analog source is to be restored, some form of jitter reduction must be performed prior to restoration, as well.

RESULTS FROM A CONCRETE EXPERIMENT WITH FRANCE TV

Production Workflow Description

France TV showcases innovations in the media field each year during the French Open tennis tournament. In 2020, the event took place in September and was dedicated to augmented tennis in various forms (7). One request of France TV was to explore the performance of state-of-art video restoration and super resolution techniques. Indeed, the COVID-19 global health crisis created a lack of new live content due to cancellation of many live events, and an increase in video consumption due to home confinement, resulting in a program supply shortage. Broadcasters and telco operators had to dig up old SD footage from their catalogs and in doing so discovered a consumer appetite for classics (i.e., legendary sport events, concerts, films or series). Therefore, France TV wanted to see how SR techniques behave in a SD to HD upscaling or even UHD 4K upscaling scenario, compared with the best-in-class filters, typically Lanczos filters, used today with real-time TV broadcasting systems.

Figure 2 highlights the 4K production workflow. It is composed of three separate workflows. On the top line, the SD or HD content directly feeds a 4K HEVC live encoder from Harmonic that deinterlaces, when needed, and upscales the content to UHD 4K resolution, using a Lanczos filter, before encoding it at a high bitrate. On the central line, the SD or HD content is processed offline with a hybrid scheme from Pixop associating a CNN-based deep-restoration (PDR) and an ML-based SR (PSR), as described in the previous section. On the bottom line, the SD or HD content is processed offline with a UNet-based deep-restoration and SR pipeline. In these last two workflows, the 4K file that is produced is then encoded at a high bitrate with the same 4K HEVC encoder as in the top line. A high bitrate is chosen for the encoding to avoid any compression impact, which would affect the comparison of the baseband pre-processing.

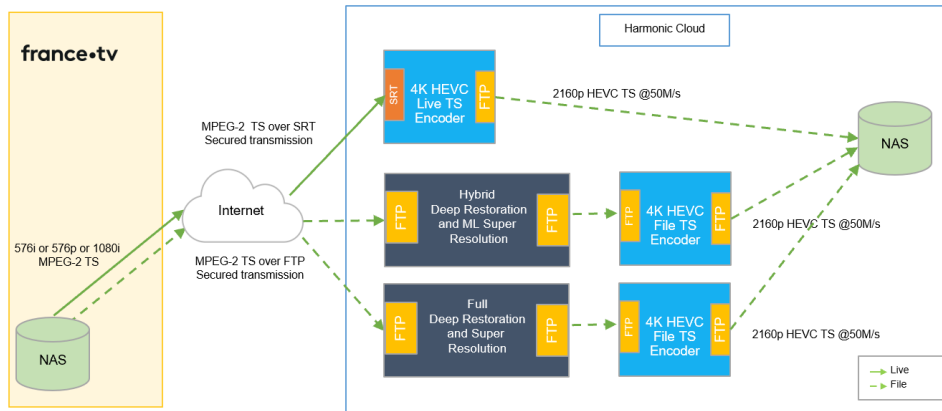


Figure 2 – 4K production workflow

Two use cases were identified for this comparison. The first use case makes use of recent SD footage with good quality, while the second use case investigates the behavior of SR on old SD footage with poor quality.

For the first use case, SD content of good quality from the 2013 French Open tennis tournament in 1024x576p@29.97Hz format was chosen. This content has been processed with the three workflows described above to compare their relative quality. In the hybrid workflow, a deep restoration is applied first on the SD content to restore texture details while upscaling to full HD (1920x1080p). The SR is then applied to generate UHD 4K content with higher sharpness and edge preservation. In the case of a full HD feed, the SR algorithm could be applied without any preliminary deep restoration phase.



Figure 3 – SR of good quality SD content

For the second use case, old SD footage from the 1984 French Open tennis tournament was chosen, as shown in Figure 4.




Figure 4 – Old SD footage



Even if the content was available in 1920x1080i@50Hz, it was originally captured in the SD format and converted in HD with poor quality showing excessive ringing, color bleedings, noise and interlaced comb effects on edges. In this second use case, a direct upscaling of the content leads to amplification of issues. Therefore, the goal is to look for deep-learning techniques that could clean the content before applying any SR or upscaling process. In this case, the processing workflow is composed of three steps:

- The first process is a denoiser, making use of deep-learning technology to “clean” the content as much as possible while keeping the useful information.
- The second process is the deep restoration, as in the first use case, to finalize the cleaning and restore details in textures while upscaling the content to HD.
- The third process is the SR to upscale to UHD 4K.

Evaluation Workflow Description

The evaluation goal is to compare the quality of the different processing techniques that can be applied in real time or offline in the headend, as described above. It may be interesting to add in the comparison of a direct upscale of the SD or HD feed in the 4K TV, as illustrated in Figure 5.

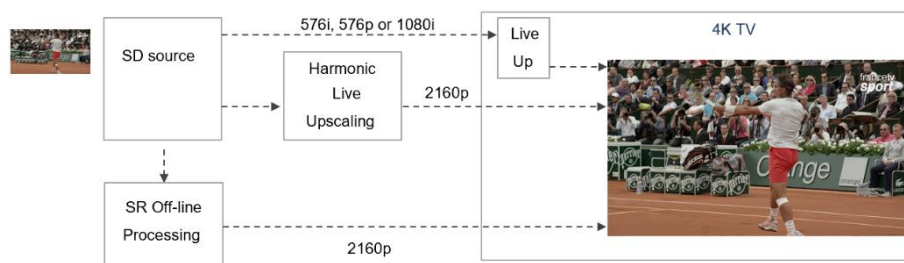


Figure 5 – Evaluation of upscaling solutions

The comparison of quality of different processing techniques has always been a challenge and is even more difficult when addressing UHD 4K resolution. Watching two solutions using a side-by-side mode, with two identical UHD TV screens, there are issues:

- It is extremely difficult to get the same rendering, even when using the same settings on two identical models.
- It is nearly impossible for both eyes to focus on the two screens at the same time given the size of UHD 4K TV screens or the evaluation distance becomes too high to be in the real-life situation.
- The use of two screens create synchronization complexities between the two.

Watching the two solutions in a sequential way offers insight on potential issues, but when talking about sharpness, there is no better evaluation than a direct relative watching.

For these reasons, the best solution is to build a split-screen mode with half the horizontal resolution on one single screen, as shown in Figure 6. Our eyes are able to watch the two sides at the same time and immediately detect which one provides the better sharpness, contrast, as well as stability.



Figure 6 – 4K split-screen mode between two solutions

Figure 7 shows the workflow for the production of split-screen streams.

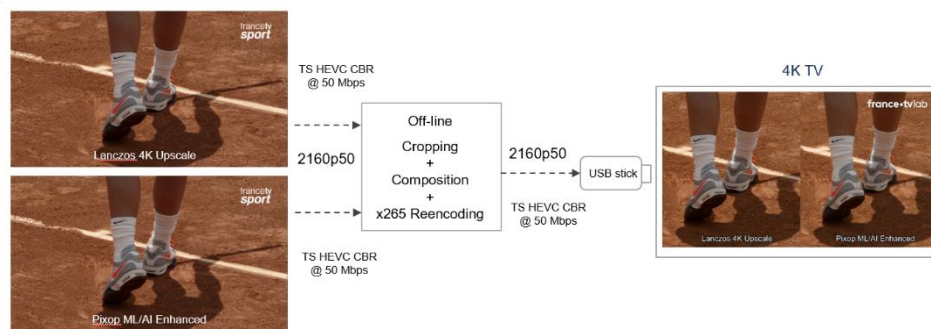


Figure 7 – Split-screen production workflow

The streams produced with live or file-based encoders at high bitrates are decoded, half horizontal part of each video is cropped in the middle, and then a side-by-side horizontal composition is made and re-encoded in HEVC at a high bitrate using a file-based encoder. Once produced, the final streams offer a demo kit that is very simple to use (i.e., play a stream on a USB stick) and is compatible with any 4K TV. Given the COVID-19 situation and home confinement, this is an immeasurable argument and the diversity of 4K TV screens that can be used provide an added value for the conclusions that can be drawn.

To simulate the upscaling made in a TV, we upscaled the SD content into UHD 4K with a bilinear filter.

The evaluation kit is composed of the following split-screen comparisons:

- Hybrid scheme vs consumer products live upscaling (bilinear filter)
- Hybrid scheme vs high-end live upscaling (Lanczos filter)
- Hybrid scheme vs full deep restoration and SR

Video Quality Evaluation Results

Video quality was assessed by video experts only on an LG OLED 65G6V 65-inch 4K TV. Viewing sessions were organized in daytime light conditions without artificial lights or external direct light pollution. The viewing distance was 1.5H, where H is equal to the height of the active part of the screen.

In the first part of the evaluation, we assessed the performance of the hybrid scheme in comparison with upscaling technologies already deployed, either in mass consumer



products (bilinear upscale) or in professional headends and high-end consumer products (Lanczos filter).

On the good quality SD content, the hybrid deep restoration and SR scheme provide a huge advantage compared with the bilinear upscaling, in terms of sharpness, contrasts and additional textural details. The SR content looks “true” native 4K, while the bilinear upscaled content looks HD.

The Lanczos upscaling has clear benefits in comparison with a bilinear upscaling, but once again, the SR content has a clear advantage in terms of sharpness, contrasts and additional details in textures.

What is extremely positive with the hybrid deep restoration and SR scheme is that at the same time it turns the content into a “true” 4K look and feel, it does not show off any weird creation, aliasing effect or temporal instabilities.

When degradation of the video source is severe, like in the older 1984 French Open SD analog tape recording, the observed benefit of hybrid denoising, deep restoration and SR over interpolation was not as profound. While some improvements in clarity were noticed in face close-ups and many artifacts or instabilities were removed in the content, long shots mostly looked smoothed in comparison. This is because the PDR model was engineered with the specific purpose of enhancing digital recordings, and there was too little information left in the denoised content. Adequately addressing analog tape recordings, which have been digitized, is a significantly bigger technical challenge, which will need to be resolved in the future. While a one-model-fits-all approach for video enhancement sounds attractive due to its simplicity, the best visual quality is produced by applying purpose-built models trained for addressing the general types of degradations found in the source video.

In the second part of the evaluation, we made a comparison of the hybrid scheme with a GAN-based deep restoration and SR technology developed by Topaz Labs and accessible as a native video enhancement application on the web (8).

The technology offers different presets, because it is extremely difficult to have one single processing that works for any kind of content, especially when the algorithm takes some risks to go one step further in the enhancement.

We have tested two presets of this technology: Gaia-HQ and Artemis-HQ, using September 2020 release. The Gaia-HQ preset seems to be more conservative than Artemis-HQ.

The comparison on the good quality SD content shows that Gaia-HQ has a very close look and feel and the same robustness. Artemis-HQ provides sharper results, is quite impressive on some scenes where it creates realistic extra details but also has issues: an aliasing effect could be perceived around clear edges, a noise effect could be seen around contours and very high frequencies generated on thin textures could result in a weird creation.

HW Requirements and Runtime Speed Evaluation

The hybrid scheme associating deep restoration (PDR) and SR (PSR) makes use of two different hardware platforms. The deep restoration runs on GPU architectures, while in the current state, the super resolution algorithm runs on pure CPU architectures. Deep restoration and SR can be combined in a single-step GPU processing in the future.

The SD restoration and upscaling to full HD is performed at 3 fps on a Nvidia GPU (GTX 1080 Ti) and 9.5 fps on AWS GPU instances (g4dn.12xlarge, 4 GPUs).



The SR HD to 4K upscaling is performed at about 10 fps on AWS CPU instances (c5.12xlarge instance, 48 vCPUs).

The GAN-based technology from Topaz Labs makes use of GPU architectures (Nvidia graphic card or Intel iGPU). On a Nvidia GTX 1080, SD to HD upscaling can be performed at 2.5 fps (same performance as PDR), and HD to 4K at around five times slower. Processing on Intel iGPU is five times slower than Nvidia GTX 1080, and processing on a pure CPU architecture is still possible but is 10 times slower than Nvidia GTX 1080.

DISCUSSION OF RESULTS AND OUTLOOK

Super Resolution and Deep Restoration in a Live Distribution System

In the previous sections, we have seen that SR works quite well with clean video content, but needs to be associated with dedicated deep restoration/noise reduction processes when the low-resolution content is of poor quality. Most professional live productions are in HD resolution, and when feeding linear channels, they are good enough to be upscaled in 4K without requiring a deep restoration/noise reduction step. The computing requirements of the Pixop ML-based SR algorithm for HD to 4K upscaling can be supported by powerful CPU servers already deployed in the cloud to achieve real-time processing at 50 or 60fps.

The SR algorithm could also be applied live on HD files that have been restored previously from SD or HD content, using offline GPU processing. These HD files can be archive footage that is low resolution and poor quality or can be today's productions using consumer products.

This hybrid live/offline approach has the advantage to possibly mix a live feed and video archives in a linear UHD 4K channel, while limiting file storage to HD resolution. It would speed up the development of UHD 4K channels, with a larger catalog of video content, since the majority of the content being archived or produced is not in 4K resolution.

Right Time for Migrating SR in the Headend

Up to now, upscaling to 4K has been deployed in end-user devices, such as UHD 4K TVs, but not on the headend side before delivery. The main reason is that UHD 4K TVs took a significant market share in the last few years, while the vast majority of linear channels throughout the world are still in SD or HD resolutions. The number of UHD 4K channels is still confidential, because of the extra bandwidth needed to deliver higher resolution. For the targeted bandwidth capacity, the delivery in HD resolution and upscaling into 4K in the TV has been judged as good enough.

Based on the evaluations above, this type of upscaling cannot offer a "true" 4K experience similar to what SR could achieve. Moreover, when the HD content suffers from compression artifacts due to bandwidth constraints, these artifacts can be amplified by the upscaling process, leading to a very disappointing 4K experience. Embedding CPU-intensive DR/SR algorithms in end-user devices may not be economically feasible or good for the planet's power consumption, but embedding these algorithms in the headend before delivery only makes sense if 4K can be delivered with a sufficient quality at a reasonable bandwidth (i.e., limited extra bandwidth compared with HD).

It won't be long before this happens because there are powerful technologies on the horizon that will enable it. There are, of course, new codecs, appearing on the market, such as AV1, as well as codecs just being standardized, such as VVC and EVC. However, it may take a few years before these codecs become mainstream. Alternative technologies can make use



of already deployed codecs, such as HEVC, and be applied in a straightforward manner to 4K delivery, with compatibility for the millions of devices already deployed. Low Complexity Enhancement Video Coding (LCEVC), which has just been standardized by MPEG, is one option with a two-layer encoding scheme.

The other option is Dynamic Resolution Encoding (DRE), which is already deployed by Netflix for VOD delivery and can be applied for live thanks to AI. These two technologies can limit the extra bandwidth of 4K delivery compared with HD delivery at a very reasonable level and make the upscaling of 4K in the headend a good option.

CONCLUSIONS

This paper has shown that super resolution techniques are now mature enough to be used in the professional video market and deliver a “true” 4K experience with digital video content captured at a lower resolution. The complexity of the algorithm under study is compatible with upscaling live HD content into 4K. For archives with a lower resolution and poor quality, it is recommended to first apply deep restoration and noise reduction processes that will preferably run offline and be trained to address specific degradations.

New encoding technologies enable 4K resolution to be delivered at a very reasonable bandwidth that justifies moving the SR process in the headend instead of in the end-user devices.

REFERENCES

1. Nasrollahi, K. and Moeslund, T. B., 2014. Super-resolution: A comprehensive survey. *Machine Vision & Applications*, 25(6), 1423-1468.
2. Liu, H. and Ruan, Z. and Zhao, P. and Dong, C. and Shang, F. and Liu, Y. and Yang L., 2020. Video super resolution based on deep learning: A comprehensive survey. arXiv:2007.12928v2 [cs.CV].
3. Romano, Y. and Isidoro, J. and Milanfar, P., 2016. RAISR: Rapid and Accurate Image Super Resolution. arXiv:1606.01299v3 [cs.CV].
4. Pixop Super Resolution, 2021: <https://www.pixop.com/filters/pixsr>
5. Thomas, C., 2019. Deep learning based super resolution, without using a GAN. Towards Data Science: <https://shorturl.at/npqrF>.
6. Pixop Deep Restoration, 2021: <https://www.pixop.com/filters/dvres>
7. Augmented Tennis at Roland-Garros 2020, Technical Aspects, <https://www.francetvlab.fr/en/posts/augmented-tennis-at-roland-garros-2020-technical-aspects>
8. Topaz Labs Video Enhancement, 2021: <https://www.topazlabs.com/video-enhance-ai>

ACKNOWLEDGEMENTS

We would like to thank the France TV team, including Vincent Nalpas, Jean-Paul Chevreux, and Yves-Marie Poirier, for giving us the opportunity to evaluate SR during a real live sports event and for providing the content processed with Topaz Labs technology.

We would also like to thank Thierry Fautier from Harmonic for his support to the SR project during FranceTV RGLabs event and after.