White Paper

Deep Learning Image Processing Technology / Upscaling Technology

Exploring the Technologies of
Canon's Neural network Image Processor and
Neural network Upscaling Tool
Using Deep Learning Image Processing Technology

Issued: November, 2023 - Initial edition
Revised: December, 2023 - Ver.01

Contents

Revisions:

1.Canon image processing technology applying deep learning

1.1 Producing "True-to-life imagery" down to the details

Scenes and situations change from one moment to the next, never to reappear or recur exactly the same way, but they can be recorded by a camera. Whether a stunning view seen for the first time or a moving occasion to cherish forever, these moments can be captured by a camera and preserved in photos. This white paper describes deep learning* image processing technology developed by Canon to enable these "True-to-life imagery".

* Deep learning
  A method of machine learning based on neural networks inspired by the human brain. Training a computer using large amounts of data enables desirable inferences and decisions based on features derived from that data.

1.2 Principles of photography

To begin, let us examine the general principles of photography using a digital camera. Figure 1-1 below illustrates the main factors affecting the image quality of an image of a subject when captured by a digital camera, and processed into a recorded image.

The light from the subject, which is illuminated by a light source such as sunlight, florescent light, LED light, or other light sources - is focused by the lens (1), to form an image of the subject and the scene. This image is ultimately formed by converting the energy of light into electrical signals by the image sensor (5).

Below we will examine the process in detail:
In this process illustrated in Fig. 1-1, the light of the subject passes through the various optical elements inside of the lens (1), to a glass filter inside of the camera body (2). This glass filter absorbs infrared (IR) and blocks ultraviolet (UV) spectrum light to eliminate invisible wavelengths of radiation, which can degrade image quality. The light of the subject then passes through an optical low-pass filter (3), which suppresses high-frequency moiré patterns, and reduces false coloring artifacts. Next the light then passes through the primary color filter array (4), which filters the light again, and assigns red, green and blue (RGB) color information to the light of the subject and scene for each individual pixel. The light then finally reaches at the light receiving substrate of the photo diodes within the image sensor (5), where the light intensity information for each color of light (Red, Green, Blue) is then assigned at each pixel of the image sensor.

At the end of this journey, the digital information of the light of the subject and scene is transmitted as an electrical signal to the image processor, which performs a variety of image processing functions. After this process, it is possible to process the image as desired by the photographer when the camera is set to capture RAW images, where various RAW image processing parameters can be adjusted as desired in the camera or by a software application.

Fig. 1-1

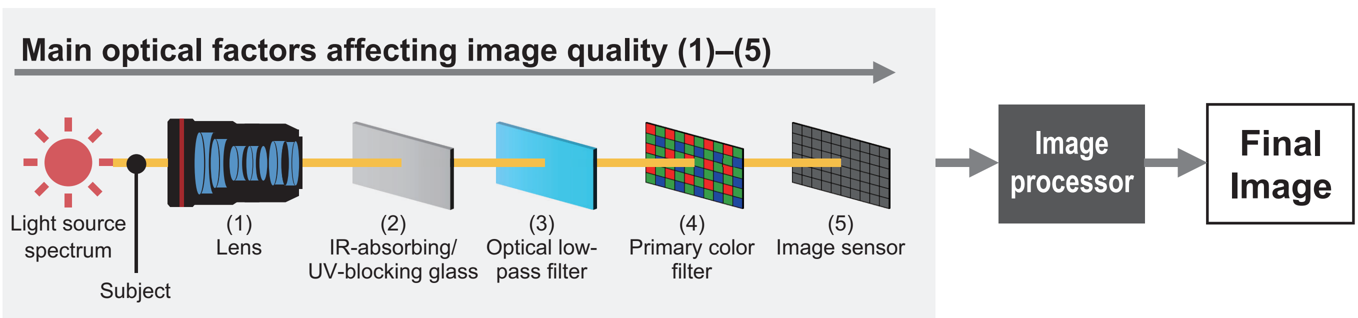

Revisions:

1.3 Inherent challenges to photographic image quality

Photography inherently is met with many challenges for obtaining optimal image quality. Examples include image noise that causes grainy photos, moiré (interference patterns), and optical aberrations inherent in lenses that cause image blur and false coloring (see Fig.1-3 to Fig.1-9). These factors affecting image quality introduce visual information that is not in the original scene.

Due to the error generated in the process of converting the energy of light into electric signals by the image sensor, a result different from the brightness and color of the original subject called noise is generated for each pixel of the image. To combat image noise, many noise reduction methods have been developed to date, but because most have negatively affected image quality in some way, by removing some subject detail during the removal of the desired noise, professional photographers are constantly looking for better solutions.

For striped or checkered pattern subjects, image sensors inherently create wavy patterns (moiré) that do not actually exist in the original subject or scene, this is due to how the image sensor's arrangement of pixels overlaps with these subjects.

Other inherent issues stem from color interpolation, which enables image sensor pixels to generate RGB data. With each pixel of the image sensor detecting only one of the three primary colors of light – red, green, or blue – detecting the other two remaining colors is inferred through a process of referring to the information of neighboring pixels known as color interpolation processing. This processing inevitably poses a risk of introducing false color (color moiré) and jagged-looking diagonal lines (jaggies) that do not really exist in the actual subject.

Although many countermeasures have been developed to preserve image quality by correcting these image quality challenges, these measures also can negatively affect the apparent resolution[*1], or affected color reproduction as a byproduct.

Still other unavoidable optical issues may arise from the lens through which cameras capture light, such as various aberrations[*2] and diffraction effects[*3] caused by lens optical characteristics. Although these aberrations are minimized by strategically using different combinations of concave, convex, and aspherical elements, in addition to special optical materials and coatings - frequently found in high-performance lenses - it is still theoretically impossible to eliminate all aberrations and their effects completely.

One example of this is the tendency toward blurring in peripheral areas of imagery created using wide-angle lenses, where optical performance is lower, when these lenses are used at small f-numbers (near maximum aperture).

*1: In this document, “Apparent resolution” refers to the sharpness of an image or image definition, and “Resolution” refers to the number of pixels in an image.

*2: Aberrations

Aberrations are a phenomena associated with the refraction of light in a lens. Types of aberration include spherical aberration (multiple points of focus caused by lens curvature), chromatic aberration (from differences in the refractive index of various wavelengths of light), and coma aberration (which causes peripheral dots that resemble comets with a tail). Astigmatism and distortion are among other optical issues, and misalignments of the imaging position which causes blurring, distortion, and color shift.

*3: Diffraction effects

Diffraction is a phenomenon that occurs when lenses and material affect the light path and disturbs it's travel. Diffraction occurs when the wave of light passes through the edge of an object in the optical path, and travels indirectly around to the shadow area of that object. This degrades image quality. Shooting at large f-numbers (with a smaller aperture) may cause light to bend around the edges of the aperture, reducing image contrast and sharpness.

1.4 Neural network Image Processing Tool

With ongoing advances in AI technology, deep learning is now used in many applications. Canon, as a company with extensive camera and lens expertise, is involved in developing original deep learning image processing to address inherent challenges to photographic image quality and producing "True-to-life imagery". In November 2022, Canon released the Neural network Image Processing Tool, an image processing software that applies deep learning in the three domains of noise reduction, color interpolation, and aberration and diffraction correction (lens blur correction).

Fig. 1-2

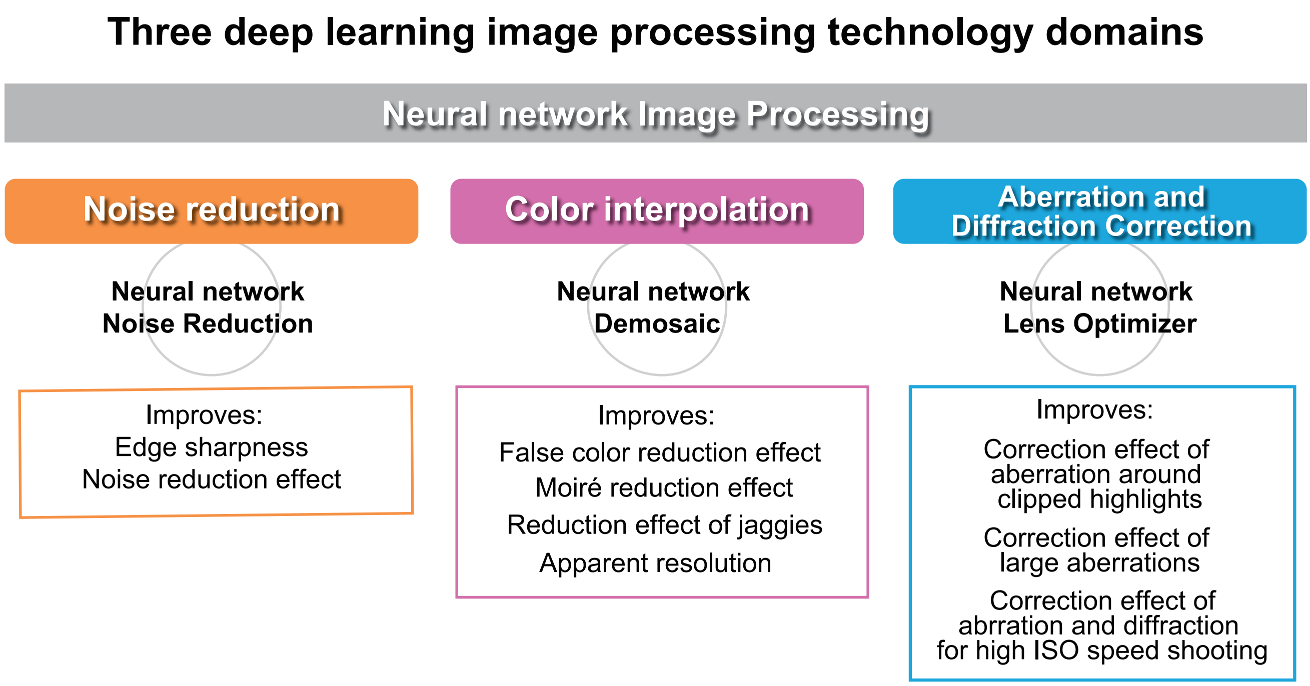

Revisions:

Noise reduction

Without a definitive solution in the domain of noise reduction in the imaging industry, Canon has pursued improvements in noise reduction image quality using image processing that applies deep learning. This technology is accomplished through training using a large amount of image data -using images with little degradation and images that are difficult to process - to produce clear, high-quality images with visible improvements.

However, deep learning approaches are not flawless, and depending on the shooting situation, they may result in “false correction” that in some cases is worse than conventional image processing.

To address this formidable challenge, Canon modified the neural network architecture itself and applied unique expertise on camera-generated image noise to refine the training process, training data, and other parameters, which led to the Neural network Noise Reduction function, which can produce clear, high-quality images.

This has made it possible to remove the noise that had been amplified along with the brightness of images during high ISO speed shooting, to obtain smooth skin tone otherwise hindered by noise, and to reproduce details and sharpen edges not easily discernible in a noisy image, among other benefits.

Fig. 1-3

Noise reduction
Reproduction of image details

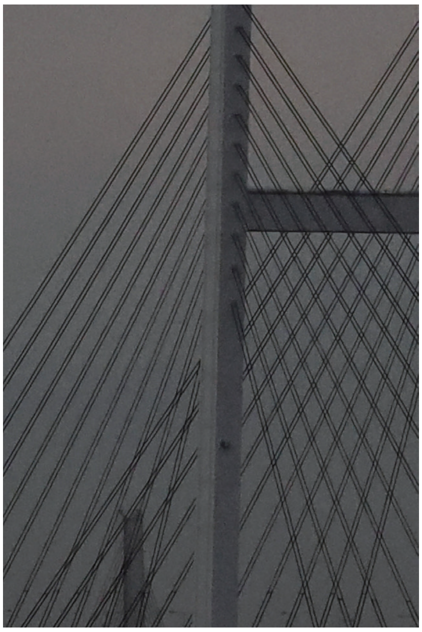

Image as captured

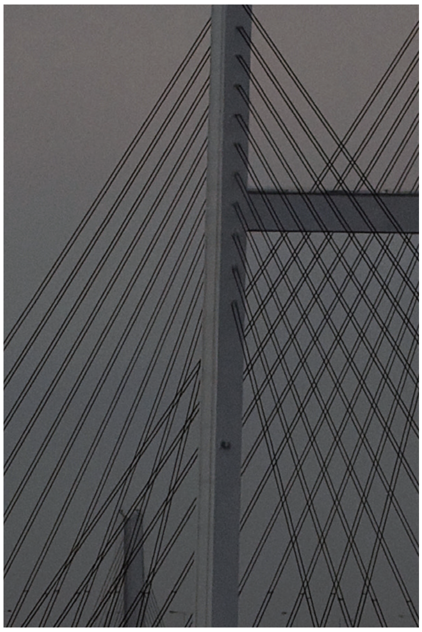

Neural network Noise Reduction

*Even at high ISO speeds, reproduction of details such as fine lines is improved.

Color interpolation

In pursuit of improvements for more accurate color interpolation, Canon has leveraged an extensive image database to establish the Neural network Demosaic function, which performs deep learning interpolation image processing. In constructing the training dataset, the developers even accounted for characteristics of human vision, which is highly sensitive to differences in brightness but less responsive to changes in color.

As a result, false interpolation has been suppressed through training focused on subjects that are difficult to infer in color interpolation. Accordingly, accurate interpolation is now possible even for challenging subjects such as stripe patterned shirts that are prone to false colors, diagonal lines that are prone to jaggies, and pet's or animal's fur/feathers that are prone to moiré or false colors, which improves apparent resolution and color reproduction.

Fig. 1-4

Color interpolation

Improved color reproduction

Image as captured | Neural network Demosaic

*Improved color interpolation also enhances color reproduction and the apparent resolution of image details.

Fig. 1-5

Color interpolation

Reduced false colors (color moiré)

Image as captured | Neural network Demosaic

*False colors in this photo of a striped shirt are greatly reduced, improving apparent resolution.

Aberration and diffraction correction

In-regards to aberration and diffraction correction, affected images caused by optical issues can be corrected and the apparent resolution can be greatly enhanced by Canon's Neural network Lens Optimizer. In developing lenses, Canon gains in-depth knowledge of the aberrations and diffraction effects inherent to each lens design.

Thus with this knowledge, design values for each lens have been applied in the photography process simulation described later (see section 1.7) to generate and learn from a large amount of deep learning training data, which has enabled correction of various blurring, such as in the peripheral areas of an image.

Additionally, the kind of correction performed by the existing Digital Lens Optimizer*[4] has been enhanced by the Neural network Lens Optimizer to isolate and correct for blurring without increasing image noise, which can be amplified in DLO correction.

* Digital Lens Optimizer

A function that enhances apparent resolution by applying lens-specific design values to eliminate loss of apparent resolution caused by two factors related to lens imaging performance – residual aberrations and the physical phenomenon of diffraction. Digital Lens Optimizer is available in Digital Photo Professional software and applies correction when developing RAW images to JPEG or TIFF format*. Digital Lens Optimizer is also built into many cameras that apply Digital Lens Optimizer effect at the time of JPEG image recording.

*Digital Lens Optimizer is available on EOS R series cameras and RF lenses, as well as on certain EOS DIGITAL SLR cameras and EF lenses.

Fig. 1-6

Aberration and diffraction correction

Improves apparent resolution

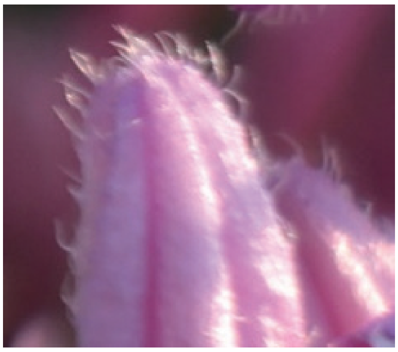

Image as captured

Neural network Lens Optimizer

*Correcting blur from optical aberrations renders subject details more sharply.

The aberration and diffraction correction is also effective in correcting blur that tends to stand out in clipped highlight subject areas, where correction is otherwise difficult due to the loss of image information from excessive brightness.

However, merely preparing data to train how to correct blurring around clipped highlights may result in unexpected false correction. Moreover, little research has been published on blurring correction related to clipped highlights. Therefore, to avoid the risk of false correction, original research by Canon was required to determine the key issues associated with the correction and to clearly understand previously unknown principles of their occurrence.

Training data was improved through computer graphics and other means, as refinements were made to the neural network architecture and image post-processing. Trial and error led to highly accurate blur correction even where blurring would normally stand out in clipped highlights.

Fig. 1-7

Aberration and diffraction correction

Blur correction of images with clipped highlights

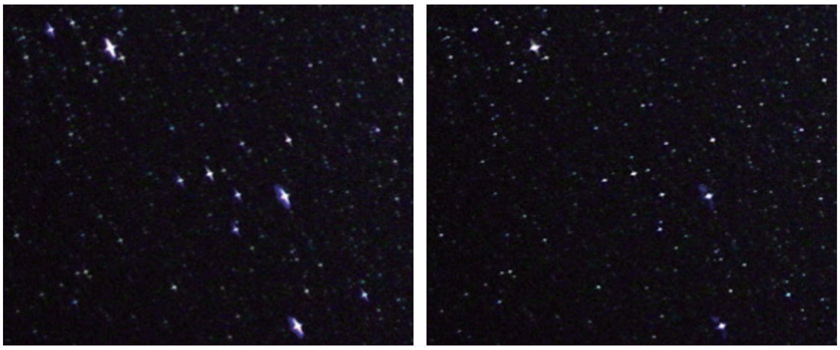

Image as captured | Neural network Lens Optimizer

*Correcting blur around clipped highlights greatly reduces peripheral blurring that tends to occur in astrophotography.

Fig. 1-8

Aberration and diffraction correction

Reduced color fringing

Image as captured | Neural network Lens Optimizer

*Color fringing around clipped highlights is also corrected, through high-resolution image processing.

Rather than applying the processing of any of these three deep learning image processing technology singularly, combining their capabilities together offers a higher level of correction, with details rendered more realistically - photos can appear more substantial and with greater definition, illustrating an overall improvement in image quality.

For example, if applying only one singular function of the total effect of the deep learning Neural network image processing such as aberration and diffraction correction (see chapter 1.4), this will only correct for blur created by aberrations and diffraction caused by the lens, however, image noise will still remain in finer image details, especially in high ISO speed images.

For this reason, in combining the additional functions of the Neural network Image Processing tool together, such as the noise reduction function in-addition to the aberration and diffraction correction function, the original image's overall image quality improvement can be further maximized saving the photographer time and effort when editing with just the use of one single software.

Fig. 1-9

1.5 Training and inference in deep learning

Deep learning image processing technology is applied in both the Neural network Image Processing Tool and the Neural network Upscaling Tool which will be described later in section 2.1.

The process of programming and training the deep learning image processing model consists of both a "Training" and a "Inference" stage, which will be discussed below. Throughout this document, and illustrated in detail in Fig. 1-10, the images corresponding to the Canon curated dataset of real-world photos taken by the camera which are candidates and would be desired to have correction performed are referred to as "Student" images. Conversely, the Canon curated dataset of ideal images which are ultimately sought as the end result and the goal of the deep learning process corrections are referred to as "Teacher" images (see 1.6 and 1.7 for details)[*1].

In the "Training" stage, the "Student" images dataset are input into the Neural network (NN), where an estimation of image correction is made through a comparison with the "Teacher" images dataset. Through this comparison, updates of the Neural Network (NN) processing parameters are generated according to the difference in the image data, thereby training the model to reduce errors and improve image quality as a process. This "Training" stage process is performed using a bulk dataset, that features a large variety of Canon curated images captured in various situations, and environments, and repeated until errors are reasonably minimal, so that corrected images estimated by the Neural network sufficiently match the "Teacher" image's image quality. Note that this "Training" stage work is conducted by Canon over the course of product development, and it is not performed on user computers.

In the "Inference" stage, the model applies it's "learned" parameters to process actual images captured by cameras. Converse to the "Training" stage, the processing for the "Inference" stage is performed on Neural network Image Processing[*2], or Neural network Upscaling Tool equipped user computers through the use of the Digital Photo Professional software application. Inference processing is applied to RAW images by the Neural network Image Processing Tool during RAW image processing (see Fig. 2-4). Whereas JPEG and TIFF images are supported for inference processing by the Neural network Upscaling Tool (see Fig. 2-1).

*1: In this document, the terms of "Teacher image" and "Student image" are used in comparison with the training data. In knowledge distillation, which is known as a method of learning a lightweight model from a large trained model, the terms "Teacher model" and "Student model" are used, but "Teacher image" and "Student image" here are different from them.

*2: Refer to the product site at the end of this document for the applicable models.

Fig. 1-10

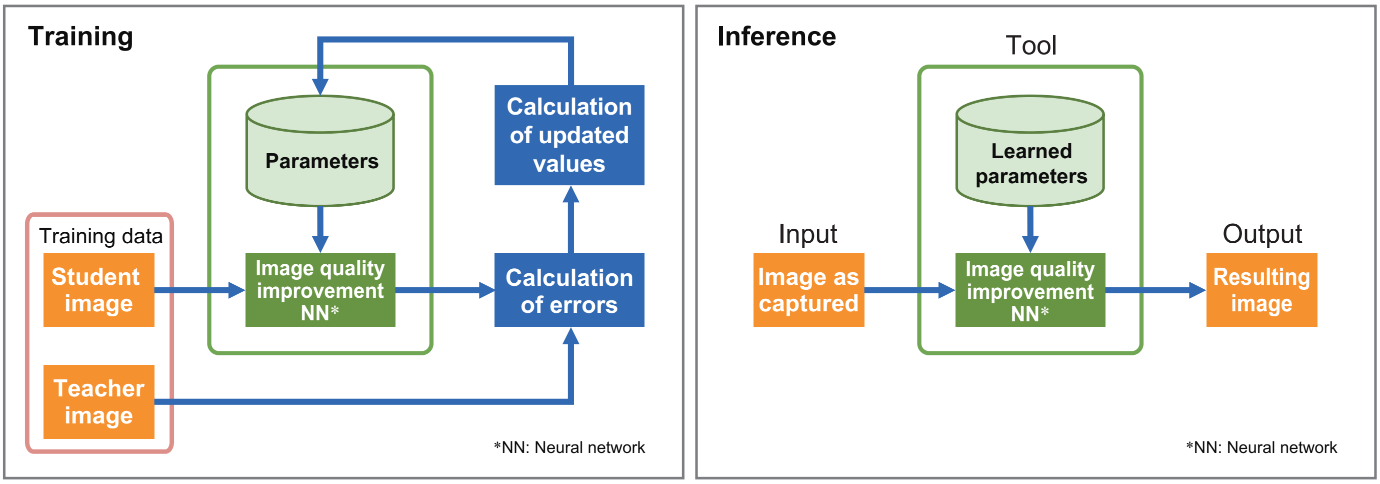

Revisions:

1.6 Training data

The key to improving the accuracy of the "Training" model's results is supplying a sufficient enough quantity of paired "Student" and "Teacher" images as Canon curated datasets, which can lead the model to learn the desired image processing goal. If an insufficient "Training" dataset is supplied, the system may not have the learned parameters to yield highly accurate results. This would occur because the "Student" image used for "Training" would not correspond to the actual real-world use image that is input during the inference stage. As a result, in less than ideal conditions in actual real-world use, the model will instead be highly dependent on actual images that closely corresponds to the limited "Training" data set. This overall condition makes the performance of deep learning systems largely dependent on large quantities of high quality "Training" data to realize the best results.

In preparing a large quantity of high-quality "Training" data, Canon draws on two main advantages as a leader in the imaging industry.

The first is the vast image database accumulated over decades of camera and lens development. A large amount of "Training" data can be generated from high-resolution RAW data captured using a camera covering an array of subjects, and this data holds more information than other formats such as JPEG.

Another Canon advantage is highly precise generation of pairs of "Teacher" and "Student" images for training from RAW data captured with use of a camera. This generation of "Training" data applies photography process simulation (see section 1.7) established through decades of experience and expertise in camera and lens development. This simulation reproduces the process of shooting in detail, reflecting the aforementioned principles of photography, for computerized simulation that generates photographic images using RAW data captured by a camera.

Together, these two Canon advantages enable generation of a vast amount of "Training" data for the development of deep learning image processing, which has led to the development of the Neural network Image Processing Tool and Neural network Upscaling Tool (see section 1-7).

1.7 Photography process simulation

The Neural network Image Processing Tool has various functions, one of which is the Neural network Lens Optimizer - a function used to correct the image quality effects of aberration and diffraction. The below discussion will detail the subject of photography process simulation to create ideal datasets for deep learning.

Photography process simulation is used to generate a pair of images as "Training" data, specifically, 1) a "Teacher" image without image blurring caused by aberration or diffraction, and 2) a "Student" image which exhibits aberration and diffraction similar to an actual image created with use of a digital camera in real world usage.

As the prospective images for the photography process simulation, the RAW data captured with a camera described earlier is used. This RAW data is processed using Canon proprietary technology to remove aberrations and diffraction-induced blurring to prepare the "Teacher" images. To prepare the "Student" images, the "Teacher" images are then carefully modified by applying the aberration and diffraction characteristics of various lenses and optical filters found in the cameras to generate RAW data modeled after actual real world photography.

Various models of cameras and interchangeable lenses can also be simulated, as can various shooting settings, which provides a large amount of ideal "Training" data. This is possible because Canon as a manufacturer oversees the entire process from design to production for its cameras and lenses, and works to maintain these design values and characteristics.

Without Canon’s proprietary technology, it would be very difficult to prepare a large, and sufficient quantity of ideal "Teacher" and "Student” images. For example, even if an image is captured in a real world shooting condition where the aberration and diffraction effects of the lens are minimal, the “Teacher” image would still retain the deterioration created during the shooting process and would be flawed compared to the images created through Canon's proprietary technology process.

Fig. 1-11

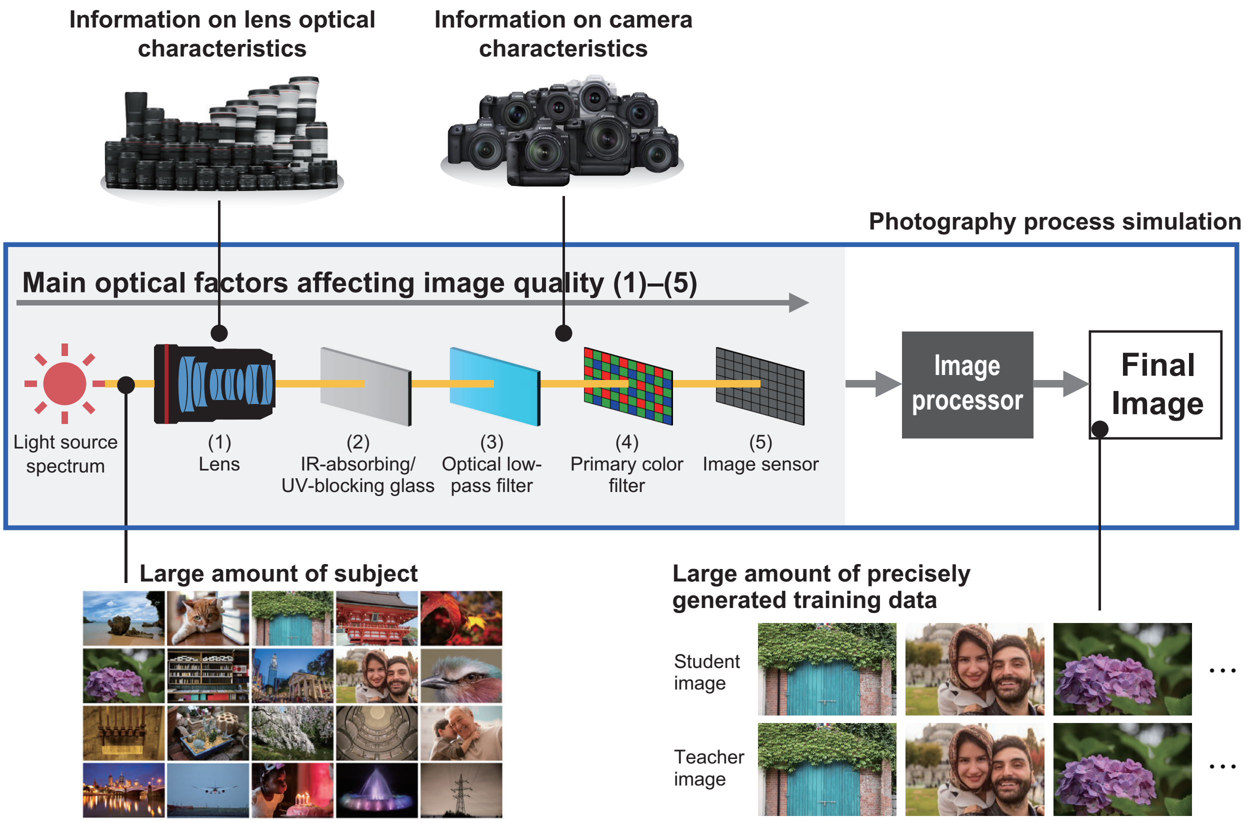

Revisions:

2. Upscaling technology

2.1 Conventional upscaling technology overview and issues

In upscaling, images are converted to a higher resolution.

The most common method of upscaling to date has been bicubic interpolation - a sharpening and enlarging technique which estimates the necessary unknown data through the brightness and color information of the known data within an image. Bicubic interpolation increases the quantity of pixels by predicting the necessary pixels to enlarge the image by sampling the color and brightness information not only from the neighboring pixels, but also the surrounding pixels to enlarge an image in a manner which makes it less likely that the subject's sharpness will be lost.

The challenge of traditional upscaling using bicubic interpolation is that as the image resolution is increased the boundary areas such as subject outlines, where white and black edges are visible, exhibit thicker outlines compared to the original image[*5]. Correcting for this effect by balancing the relationship between the thickness of the outlines and the apparent resolution conversely decreases the apparent resolution compared to the original image. Upscaled photos might also seem slightly blurrier to many photographers who will crop their image for stronger composition, and also desire to enlarge the remaining pixels for print enlargement. Examples of this often include wildlife, or sports photography, or even city scenes captured while traveling.

*5: For details, see “3.3 Comparison with conventional upscaling technology.”

2.2 Neural network Upscaling Tool overview

In April 2023, Canon released the Neural network Upscaling Tool, image processing software that applies deep learning technology. The software makes it possible to double the number of vertical and horizontal pixels and quadruple the total number of pixels while maintaining the apparent resolution of the source JPEG or TIFF image.

Apparent resolution is sometimes lost in conventional upscaling when images are enlarged or cropped (see Fig. 2-1, 2-4). In contrast, the Neural network Upscaling Tool applies deep learning image processing technology to generate images near the apparent resolution of the original subject without altering color, brightness, or noise.

2.3 Benefits and uses of Neural network Upscaling Tool

One use of the Neural network Upscaling Tool is in preparing images with a higher resolution for large-format printing. Another applies to images that have a lower resolution after cropping, for which users may wish to maintain sharpness by increasing the resolution.

These upscaled images are as life-like as the original image while minimizing blur at the subject boundaries and additionally maintaining the natural sense of sharpness that is visible in the original image size. The effect of upscaling while maintaining high apparent resolution is especially noticeable in fine animal fur, buildings or text with defined edges, or expansive wide-angle landscape photos.

*Note: The tool can also be used with images captured with Canon cameras already on the market and cameras of other manufacturers.

Revisions:

Fig. 2-1

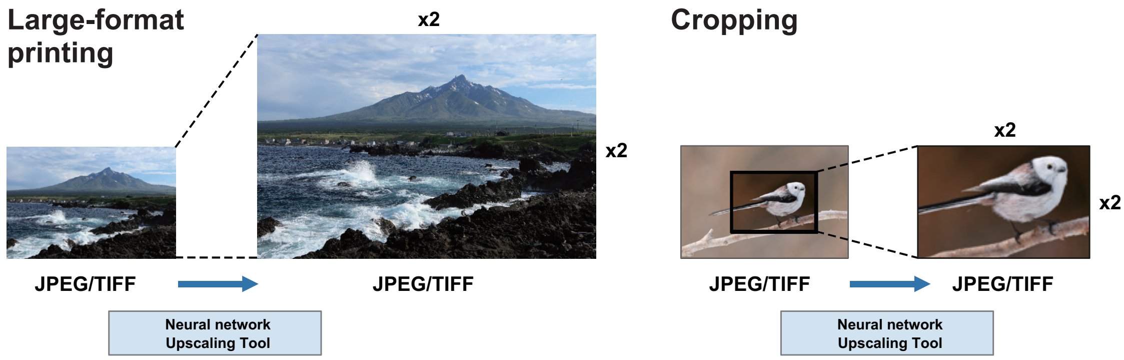

Fig. 2-2

Typical processing (bicubic interpolation)

Canon upscaling

Fig. 2-3

Typical processing (bicubic interpolation)

Canon upscaling

The full benefit of deep learning image processing is available when using the Neural network Upscaling Tool to upscale JPEG or TIFF images saved after processing with the Neural network Image Processing Tool in Digital Photo Professional as described previously.

Fig. 2-4

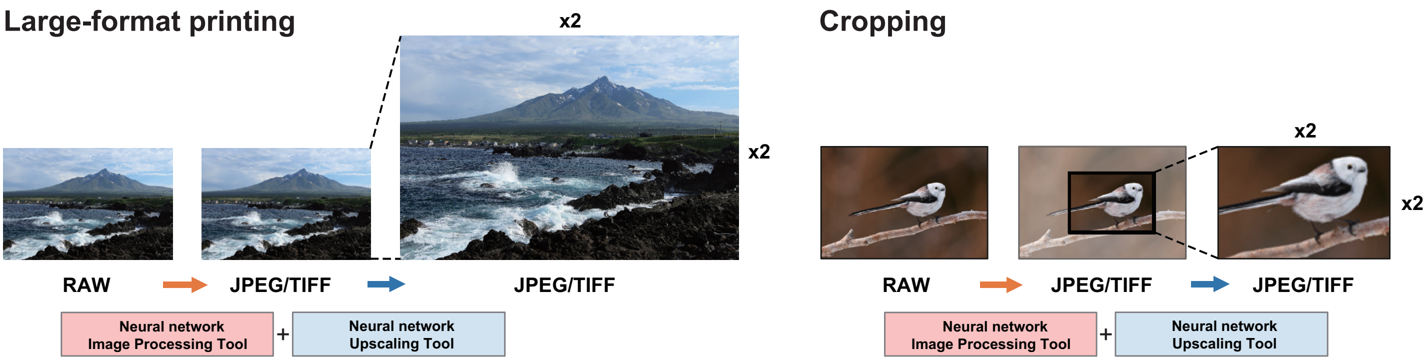

Fig. 2-5

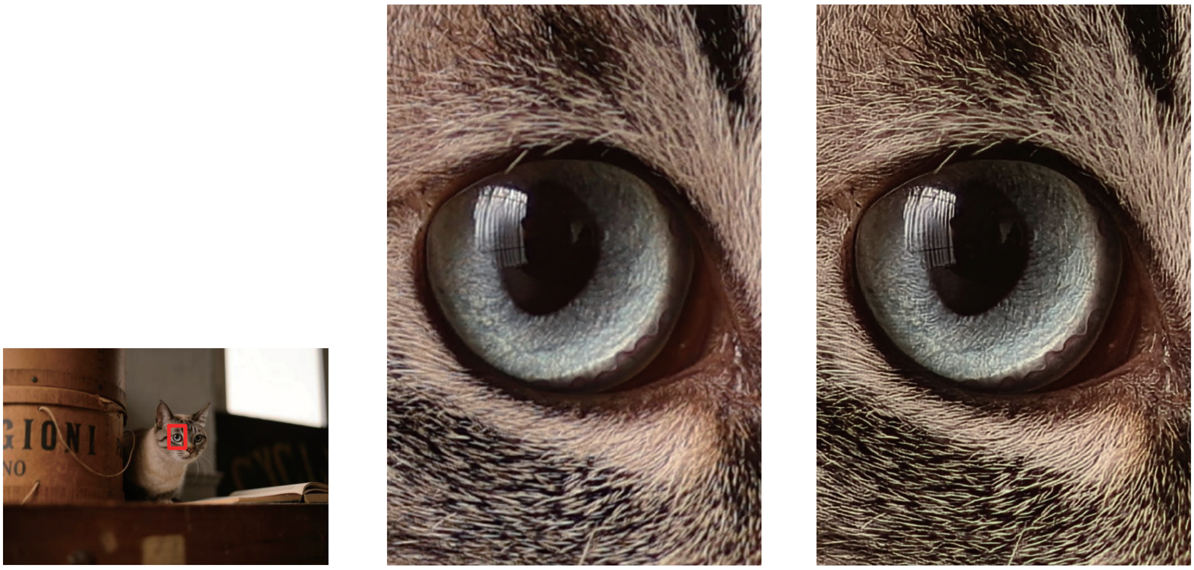

Typical processing (bicubic interpolation)

Canon deep learning

Revisions:

3. Features and advantages of Canon upscaling technology

3.1 Target image quality

As pursued by Canon, this upscaling is intended to maintain the original apparent resolution when photos are enlarged. Thus, captured images are processed so that an image enlarged by deep learning image processing nearly matches the original high resolution of the image. Although excessive enhancement of resolution beyond the original image resolution is technically possible with deep learning, the results would diverge from the impression given by the original photo. The Neural network Upscaling Tool was designed paying close attention to details such as color, brightness, and noise to maintain the impression of the original images provenance even after enlargement.

3.2 Obtaining the target image quality

"Training" and "Inference" to ensure target image quality are as follows.
In the "Training" stage, the Neural network Upscaling Tool employs "Training" data that is generated very precisely using the same photography simulation discussed for the Neural network Image Processing Tool. This photography process simulation is used to generate a pair of images as "Training" data, specifically, a high-resolution "Teacher" image corresponding to the enlarged image and a lower-resolution "Student" image corresponding to an actual image captured by users in the real world.

As source images for simulation, the RAW data described previously is used. "Teacher" images are obtained by editing this RAW data. "Teacher" images are then used to create lower-resolution images with characteristics of various optical filters found in the cameras applied, enabling generation of "Student" images modeled after actual photography.

In this way, training that employs "Teacher" and "Student" images which faithfully reproduce the relationship between the image as captured, and the image as enlarged makes it possible to modify an image so that it retains nearly the original resolution after enlargement, instead of generating high-resolution images not based on reality.
In the "Inference" stage, user images are supplied to a trained neural network to produce enlarged versions. Although as mentioned, the neural network is trained in a way faithful to reality, close examination of image quality may reveal slight over-correction for some scenes.

As a company with high standards in imaging quality, Canon has been committed to obtaining the target image quality as reflected in details such as color, brightness, and noise of the resulting images. To this end, another step of image processing has been added to adjust image quality so that deep learning image processing can be used with confidence.

Fig. 3-1

Large amount of precisely generated training data for upscaling

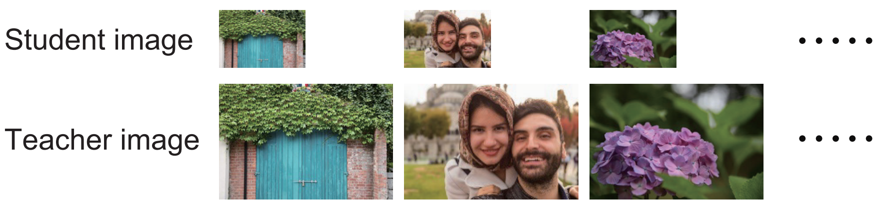

Fig. 3-2

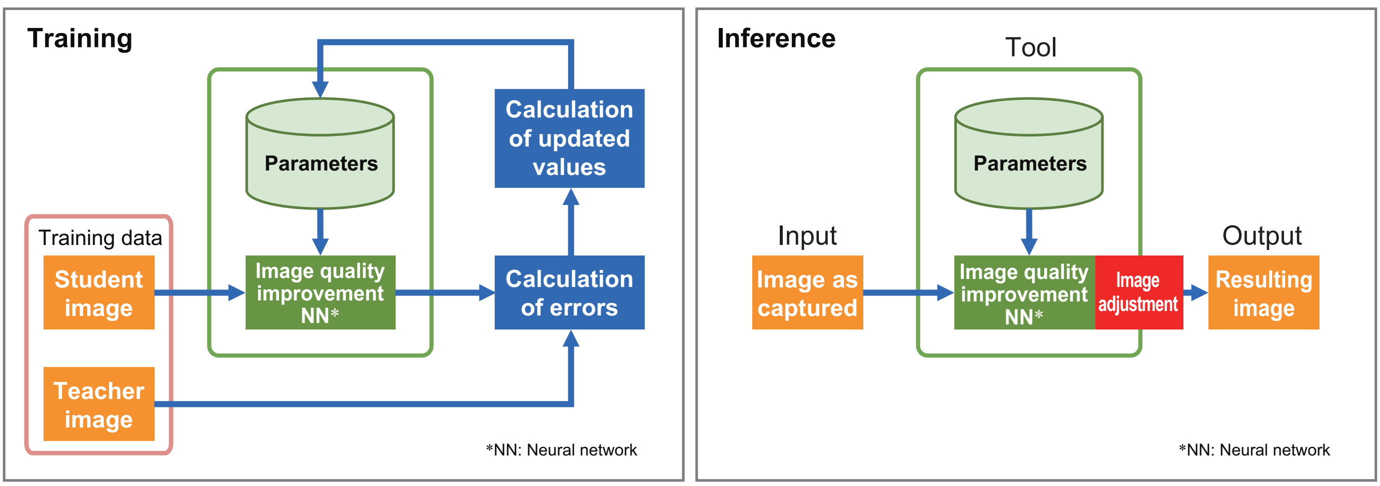

Revisions:

3.3 Comparison with conventional upscaling technology

Here, the effect of upscaling with deep learning image processing is illustrated with a schematic diagram in the case of 2× enlargement.

Charts (a)–(d) on the right in the following figure show the brightness (line profile) per pixel in a cross-section view of the blue-line portion (with the length indicated by “▬” enlarged) of the black-and-white edges of images (A)–(D) at left in the figure. In charts (a)–(d), pixel positions (coordinates) are represented horizontally by dots, and vertically, brightness increases toward the top and decreases toward the bottom.

Chart (a) is the line profile of the Original image before enlargement, showing that the left side of the blue line is darker and the right side is brighter. Upscaling with the conventional enlargement method of Bilinear interpolation, shown in chart (b), involves stretching the Original image to double its horizontal length and interpolating the pixels required in between from the brightness of neighboring pixels (●) to generate new pixels (●). This doubles the pixel count but changes the contrast ratio of the line brightness gradient when compared to the Original image, resulting in blurry edges.

Another conventional scaling method is bicubic interpolation, which includes sharpening to maintain apparent resolution. As shown in chart (c), although the line brightness gradient of Bicubic interpolation is improved when compared to (b) Bilinear interpolation, and the image more closely resembles the original, pixels at both ends of the edges where the black and white pixels meet (●) are over-corrected.

In contrast, as shown in chart (d), Canon Upscaling processing applies deep learning to estimate pixels (●) and brings the line brightness gradient closer to the original image while maintaining brightness by the edges, producing a clear upscaled image with no loss of original image quality. As a result, compared to conventional upscaling, subject edges can be sharper, hair finer, and intricate structures clearer, with a greater sense of texture.

*Note: Scaling and sharpening methods described here are only examples, and several other methods are available.

Fig. 3-3

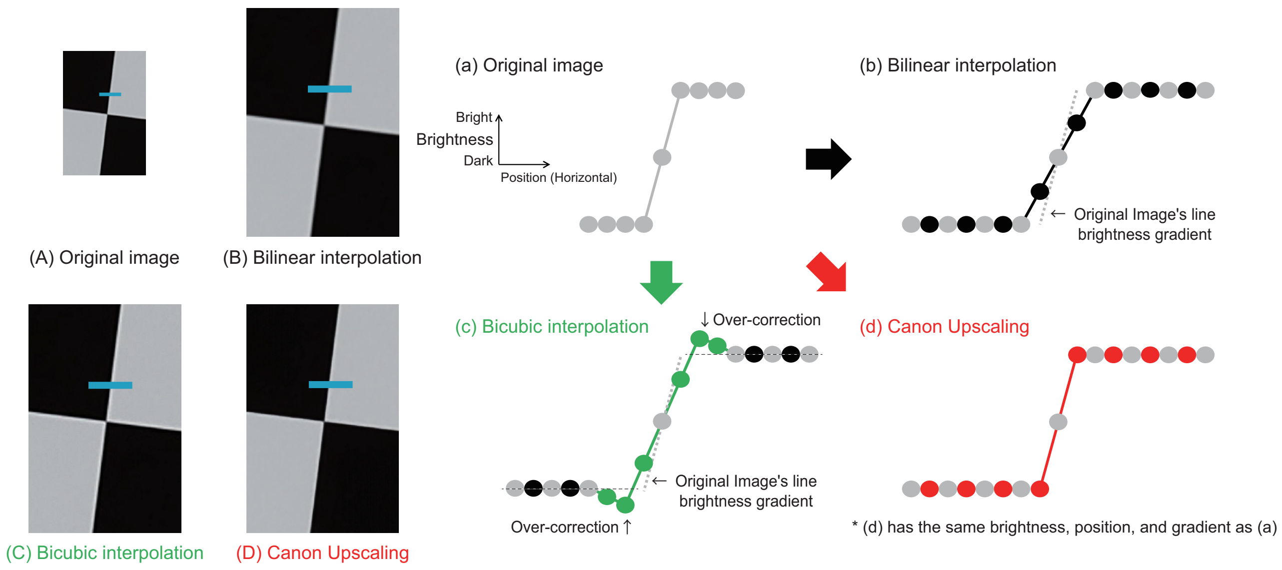

4. Image quality of Canon pursuing “True-to-life imagery”

The Neural network Upscaling Tool was developed to enlarge images without altering color, brightness, or noise, to enable “True-to-life imagery” down to the details. In the "Training" stage, a large quantity of "Teacher" and "Student" image pairs are generated, with the former having double the pixel count vertically and horizontally through photography simulation faithful to actual shooting, as described. To learn the relationship between the pairs of low resolution "Student" images, and the high resolution "Teacher" images (Fig 4-1), a large amount of data set images are required in the "Training" stage.
Specifically, lower-resolution "Student" images include traces of fine subject structures that could not be expressed by the available pixels, in the form of an aliased signal. This provides the means for determining relationships (characteristics in common) between low- and high-resolution images. Aliasing is a phenomenon which in principle occurs in signal sampling. A subject structure finer than the sensor pixels appears as a structure rougher than the pixels.
For example, as in the low-resolution image on the left in the following figure, sharp edges appear as “jaggies,” forming patterns in photos that do not actually exist. When a user’s captured image is upscaled with deep learning image processing, the neural network detects the trace aliasing found in fine details of the source image and estimates the pixels needed for a high-resolution image to produce the enlarged image. In summery, the learnings from the "Training" stage are designed to avoid rendering (generating) false subject structures where subjects are actually smooth or blank.

Fig. 4-1

Low-resolution image | High-resolution image

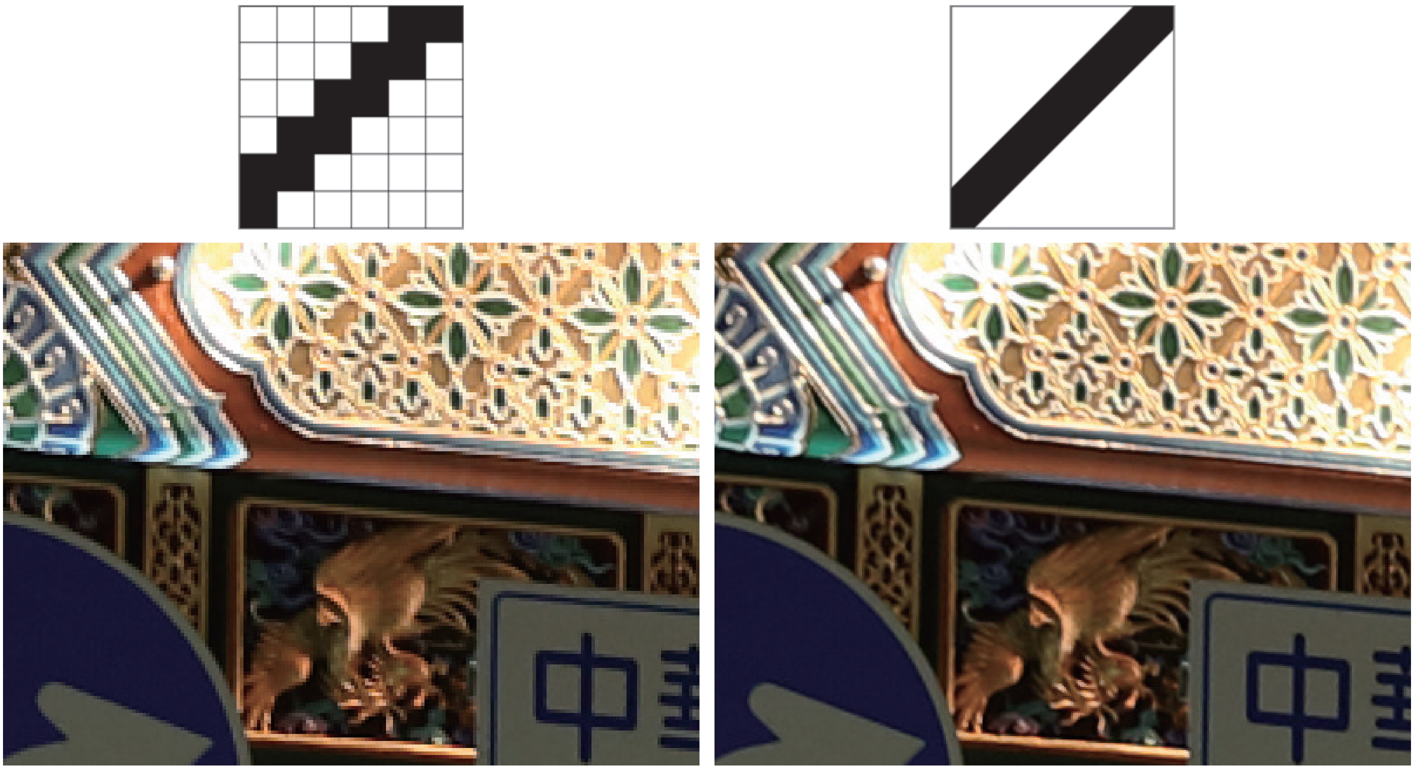

Revisions:

Moreover, recent years have seen the release and growing popularity of AI technology that uses diffusion models or similar, which to generate wholly new images in response to text descriptions of characteristics of a desired image. For example, AI technology can generate fine details such as eyelashes or detailed facial wrinkles on a subject's face in a portrait that is generated from a low resolution source image that originally would be so blurry that facial features are difficult to discern.

In this context, we can imagine an image of a face that happens to be available as data with such low resolution that the person cannot be identified. Countless versions of similar high-resolution faces exist that would become identical to this low-resolution image when the resolution is reduced. Even if a low-resolution face includes no traces of subtle facial features such as eyelashes or wrinkles, an AI can recognize that it is an image of a face and can generate a similar version within its scope of training. Thus, there is no assurance whatsoever that the high-resolution faces created will match the true facial details of the original low resolution face. This applies not only to increasing the resolution of faces but also to sharpening images that are so blurry that the structure of the subject cannot be discerned, for example, let us examine the comparison below Fig. 4-2.

The image of the overall scene (A) was captured with the rightmost wine bottle in focus, and each subsequent bottle to its left is intentionally and increasingly out of focus and blurry the farther that it is from the focus plane, and the depth of field (DOF). The images below (B to E) it correspond to the out-of-focus area outlined in red in the overall photo (A). Therefore, the image (D) upscaled through other processing methods exhibits higher, but undesired image resolution improvements to the point that the expressiveness of the bokeh is lost, and an image different than the actual bottle label is generated losing the original photographic intent. In contrast, Canon deep learning image processing stands apart from these approaches (C). As mentioned, image processing reflects the photography simulation and is based on relationships between image elements at different resolutions, which produces natural results without losing the expressiveness of bokeh.

Fig. 4-2

A: Overall photo
*Area outlined in red is blurry because it is out of the depth of field.

B: Conventional processing
(Bicubic interpolation)

C: Canon deep learning

D: Example of other method

E: For reference: Image in focus

As a leading manufacturer in the imaging industry, Canon has been developing products that enable Canon photographic image quality – "True-to-life imagery" down to the details. We will continue to pursue technological advances and a satisfying photography experience for Canon camera and lens users.

[Reference]

Deep Learning Image Processing Technology
https://global.canon/en/technology/dl-iptechnology-2023.html

Upscaling Technology
https://global.canon/en/technology/dl-upscaling-2023.html

Neural network Image Processing Tool
https://sas.image.canon/st/nnip.html

Neural network Upscaling Tool
https://sas.image.canon/st/nnups.html

* The "Deep Learning Image Processing Technology/Upscaling Technology" described in this document provides an overview of our company's technology as of November 2023. Therefore, the contents of this document are not guaranteed in the future.
* Effects and results may differ from those described in this document depending on the subject, shooting conditions, etc.

Revisions: