

# Denoising Diffusion Probabilistic Model Inspired Deep Learning for Single Image Super Resolution

Sophia Tesic

Liam Cawley

Alexandra Lavacek

Namitha John

Dennis Lin

## Abstract

*This paper will cover our processes and exploration of solutions for Super Resolution. Specifically, our goal is to explore complete solutions and various components of related works that have the potential of beating current super resolution baselines.*

## 1. Introduction

Super resolution is an area of increasing interest in computer vision with the general idea being to produce a high resolution image from a low resolution input or set of inputs. Super resolution is used frequently in medical imaging, satellites, security surveillance, and more, as implementing super resolution is generally inexpensive and as a result, can help alleviate issues with lack of high resolution output from existing hardware, illustrating the importance of super resolution as a whole. To explore improving existing super resolution techniques, we began with a literature review on various proposed solutions for efficiently creating higher resolution images. A selection of papers we read explored Bayesian super resolution, super resolution with inpainting, and super resolution using deep CNNs, among other solutions. Current efforts in super resolution successfully complete the task of creating high resolution images, but many models' outputs include artifacts and blurred details, or are overly smooth. Therefore, our problem statement is that there is a need for models that produce high resolution output that creates desired detail without undesired noise. After exploratory research into deep learning approaches to super resolution, we decided to develop an enhanced model inspired by Denoising Diffusion Probabilistic Models to address single image super resolution. While DDPMs show promising results for image generation tasks, we introduce several unique features and modifications that adapt to the super resolution domain. While a normal DDPM is typically used for image generation tasks, by learning a reverse diffusion process and denoising images by estimating the noise added at each step and iteratively refining the image, our advanced model is specifi-

cally designed for the task of single image super-resolution. Instead of generating new images from noise, the model focuses on enhancing the resolution and quality of existing low resolution images. The model learns to map low resolution images to their high resolution counterparts directly. In terms of architecture, it is typical of DDPMs to employ a U-Net like architecture, where the encoder gradually down-samples the input image and the decoder up-samples the latent representation to generate the output image. The encoder and decoder are connected through skip connections to preserve spatial information. Our model also adopts a U-Net like structure, but through the incorporation of residual blocks in both the encoder and decoder, we facilitate the effective learning of deeper connections. These residual connections also help alleviate the vanishing gradient problem and enable the model to capture complex patterns.

## 2. Background

As previously mentioned, we first began with exploring current proposed solutions to super resolution. The baseline solution that we explored, bicubic interpolation, does not use deep learning. However, we were interested in exploring a solution that was based on deep learning, as we believed it would help us surpass baseline performance.


The primary works that we used were *Image Super-Resolution Using Deep Convolutional Networks* by Chao Dong and Chen Change Loy. We also explored *Super Resolution using Edge Prior and Single Image Detail Synthesis* by Yu-Wing Tai, Shuaicheng Liu<sup>2</sup>, Michael S. Brown, and Stephen Lin.

To follow the methodology of our method, the reader will need a basic understanding of super resolution and the common technique of bicubic interpolation. The reader will also need a basic understanding of common computer vision techniques in deep learning and traditional computer vision. Below is a basic explanation of the topics we hope our readers understand prior to exploring our solution. We also provide a list of key terms which readers may find useful to independently research in

order to better understand the content of this paper, particularly the following methods section, and adjoining code.

## 2.1. What is Super Resolution?

Super resolution refers to enhancing the resolution of an image, typically from a lower resolution version to a higher resolution version. This enhancement is done in a way that aims to preserve or even improve the sharpness and quality of the image. It's used in various fields such as image processing, computer vision, and photography to improve the visual fidelity of images.



## 2.2. What is Bicubic interpolation with regards to Super Resolution?

Bicubic interpolation is a method often used in super resolution to upscale images. It works by estimating pixel values based on surrounding pixels, producing smoother results compared to simpler interpolation methods like bilinear. Additionally, this method helps reduce discrepancies and improve the overall quality of the upscaled image in super resolution applications.

## 2.3. What is Deep Learning?

Deep learning is a branch of machine learning that focuses on training artificial neural networks to learn from large amounts of data. In the context of images, deep learning involves using these neural networks to extract features from pixel data. By iteratively adjusting the network's parameters through backpropagation, deep learning algorithms can progressively improve their ability to understand and interpret complex visual information, ultimately enabling a wide range of image-related applications.

## 2.4. Key Terms

Convolutional Neural Network, U-Net, encoder, decoder, activation function, perceptual loss, skip connection, hyperparameter tuning

## 3. Methodology

We will first explain several methods which we researched. Then, we will outline our final solution.

### 3.1. Exploratory methodology

The first deep learning solution we explored was super resolution with image repaint. There are a few solutions to super resolution in existing literature that use image repainting, but one of the more compelling options we saw was strikingly similar to how we accomplished homework number five's question six. The solution went as follows:

1. a low-resolution image is first built from the original picture;
2. an inpainting algorithm is applied to fill-in the holes of the low-resolution picture;
3. the quality of the inpainted regions is improved by using a single-image SR method" (Meur & Guillemot, 2012).

However, one problem we knew we could face is solving the inverse problem of step 2, or rather figuring out the mask needed to fill-in the holes of the low-resolution picture. With this challenge, we decided to table the Image Repaint solution.

We moved on to explore image super-resolution using deep convolution networks (CNNs). We explored solutions that allowed for an "end-to-end mapping between low-high resolution images" (Dong & Loy, 2016). The model we primarily researched, Super-Resolution Convolutional Neural Network, or SRCNN, was composed of three layers. Specifically, the first layer focuses on detecting edges and textures (key features) for each patch, while the second layer detects intensity differences. Finally, the third layer is used for reconstruction. The model we explored had a performance which could be adjusted by changing the number of filters and their sizes, but like all CNNs, this affects the trade-off between accuracy and computation time needed.

After training, compared to other super-resolution methods like bicubic interpolation, SRCNN typically outperforms on quantitative metrics such as PSNR, a measure of image quality, and SSIM, a measure of structural similarity, producing higher-quality images (Dong & Loy, 2016). Based on our research, we thought that SRCNN seemed promising, and decided that our final solution would also use CNNs to map between low-high resolution image pairs.

Additionally, we explored Enhanced Deep Residual Networks for Single Image Super-Resolution (EDSR). The "Enhanced" aspect in EDSR comes from the improvements

made over the original deep residual network (ResNet) architecture. Residual networks utilize skip connections or shortcuts to jump over some layers, allowing the network to learn residual functions easier. EDSR builds upon this concept, employing a deeper network with residual blocks to learn the mapping from low-resolution to high-resolution images more effectively (Lim et al., 2017).

## 3.2. Experimental Methodology

### 3.2.1 Baseline method from literature:

We utilize the DIV2K dataset, which is commonly used for super resolution model work (Timofte et al., n.d.). DIV2K contains 800 low resolution images, created using a downsampling factor of 2, and 800 high resolution images for training, and 100 low resolution images and 100 high resolution images for testing (Timofte et al., n.d.). For computational simplicity, we forgo use of the validation set provided in DIV2K, instead utilizing the training and testing sets. Before implementing our deep learning model however, the first step of our process, after choosing a dataset, was to implement a baseline to compare our own implementations to. To do so, we conducted a literature review of traditional algorithms for super resolution and decided to use bicubic interpolation as our baseline, as it is computationally efficient compared to other methods.

We performed all coding in python using Google Colab (see appendix). For our baseline results, we implemented bicubic interpolation by combining OpenCV’s implementation and prior work from Martin Krasser’s github repository, which provided us with a simple architecture for loading data in a way that extended to our future deep learning implementation (Krasser, 2018).

For appropriate comparable results for our deep learning model, we used images the deep learning model would be evaluated on as data for the bicubic interpolation. Therefore, we used the 100 low resolution images from DIV2K’s test set as input for bicubic interpolation and compared output to the 100 high resolution images from the test set. We evaluated the performance of bicubic interpolation, which is not a deep learning method, using peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM). The generated super-resolved images were compared against the ground truth high-resolution images. The SSIM calculation used a window size of 11 to capture local structural similarities.

### 3.2.2 Our method

To surpass the performance of bicubic interpolation for super resolution on the DIV2K dataset, we implemented two

deep learning models: a basic model and an advanced model inspired by the Denoising Diffusion Probabilistic Models (DDPM) architecture. While not strictly a DDPM, the advanced model incorporates key ideas from DDPMs and other generative models to learn the mapping between low-resolution and high-resolution image pairs, with the goal of generating high-quality, detailed output images. Python libraries we utilized for these models are datasets for DIV2K loading, OpenCV/cv2 for bicubic interpolation, pillow for basic image processing, matplotlib pyplot for image display, scikit-image for evaluation metric formulas, skimage for metric calculations, numpy for matrix operations, TensorFlow for deep learning model architecture functions and os and sys for basic python structures.

### 3.2.3 Our method: Model 1

Model 1 consists of a simple encoder-decoder structure without residual blocks or skip connections. The encoder uses convolutional layers to downsample the input image and increase the number of feature maps, while the decoder uses transposed convolutional layers to upsample the encoded features and reconstruct the high-resolution output image. This basic architecture is commonly used in image-to-image translation tasks (Isola et al., 2017). Model 1 uses a simpler, shallower architecture that might not effectively capture complex patterns or high-level features necessary for more nuanced super-resolution tasks. It relies solely on mean squared error (MSE) loss, which focuses on pixel-level accuracy. While this approach can yield decent results, it may not be sufficient to capture and reconstruct fine details and textures in the super-resolved images. Despite its limitations, Model 1 serves as a good starting point for exploring deep learning-based super resolution. It provides a baseline for comparison and helps us understand the fundamental components of a super resolution model. By analyzing the performance and shortcomings of Model 1, we can identify areas for improvement and develop more advanced architectures that incorporate techniques like residual blocks, skip connections, and perceptual loss. Model 1 was trained on the DIV2K dataset, which consists of high-quality images suitable for super resolution tasks. The dataset was split into training and validation sets, with the validation set used for evaluating the model’s performance. During training, the model was optimized using the Adam optimizer with an appropriate learning rate schedule. The model was trained for a specified number of epochs with a batch size of 1. To evaluate the performance of Model 1, we calculated average Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM) scores on the validation set. These metrics provide a quantitative measure of the quality of the super-resolved images compared to the ground truth high-resolution images. Additionally, we visu-

ally inspected the generated images to assess their perceptual quality and the presence of artifacts or blurriness. By implementing and evaluating Model 1, we establish a foundation for further exploration and improvement in our super resolution project. The insights gained from this model will guide our development of more advanced architectures and techniques to achieve higher-quality super-resolved images. In appendix, see `base.py` for the code implementation of Model 1.

### 3.2.4 Our method: Model 2

Building upon Model 1, we developed an enhanced model that incorporates several improvements to enhance the quality of the super-resolved images. Model 2 is designed to address the limitations of Model 1 and leverage techniques that have proven effective in recent super resolution research. One of the key additions to Model 2 is the use of residual blocks (He et al., 2016). Residual blocks have been widely adopted in deep learning architectures, particularly in the field of super resolution, due to their ability to facilitate the learning of deep representations. In our enhanced model, we include six residual blocks, each consisting of two convolutional layers with  $3 \times 3$  filters, followed by LeakyReLU activation. The residual connection allows the network to learn residual functions effectively, mitigating the vanishing gradient problem and enabling the flow of information across multiple layers. Model 1’s performance was limited by its shallow architecture, which may not capture complex patterns and high-level features necessary for nuanced super resolution. By incorporating residual blocks, Model 2 can learn more sophisticated representations and generate higher-quality super-resolved images. The residual blocks also facilitate the training of deeper networks, allowing the model to capture a wider range of image features and details. Another significant enhancement in Model 2 is the introduction of channel attention (Hu et al., 2018). Channel attention allows the model to adaptively weigh the importance of different feature channels based on their relevance to the super resolution task. In our implementation, we employ a channel attention mechanism that learns to assign weights to each channel using a squeeze-and-excitation block. The squeeze operation globally averages the feature maps, reducing their spatial dimensions, while the excitation operation learns channel-wise weights through a series of fully connected layers. By applying channel attention after each residual block, the model can prioritize informative features and suppress less relevant ones, leading to improved super resolution performance. The decision to incorporate channel attention was motivated by the observation that not all feature channels contribute equally to the super resolution process. Some channels may capture more relevant details and textures, while others may contain less

useful information. By allowing the model to adaptively focus on the most informative channels, channel attention enhances the model’s ability to reconstruct fine details and produce visually pleasing results. In addition to residual blocks and channel attention, Model 2 also employs a larger kernel size of  $5 \times 5$  in the first and last convolutional layers. This design choice allows the model to capture a wider receptive field and consider more contextual information when processing the input and generating the output. The larger kernel size enables the model to better handle complex textures and structures in the super resolution task. Furthermore, Model 2 incorporates a global residual connection, where the upsampled input image is added to the output of the network. This global residual learning strategy helps the model focus on learning the high-frequency details and residuals necessary for accurate super resolution. By directly passing the upsampled input to the output, the model can concentrate on refining the details and correcting any artifacts introduced during the super resolution process. The training procedure for Model 2 follows a similar approach to Model 1. The DIV2K dataset is used for training and validation, with the model being optimized using the Adam optimizer and a mean absolute error loss function. The model is trained for 100 epochs, and the batch size is set to 1 to accommodate memory constraints. To evaluate the performance of Model 2, we calculate average PSNR and SSIM scores on the validation set. These metrics provide a quantitative assessment of the quality of the super-resolved images compared to the ground truth high-resolution images. Additionally, we perform a qualitative evaluation by visually inspecting the generated images to assess their perceptual quality, sharpness, and absence of artifacts. By incorporating residual blocks, channel attention, larger kernel sizes, and a global residual connection, Model 2 aims to overcome the limitations of Model 1 and achieve state-of-the-art performance in single image super resolution. The design choices are inspired by recent advancements in deep learning-based super resolution and are tailored to capture complex image features, prioritize informative channels, and generate high-quality super-resolved images. In appendix, see `resblock_attention.py` for the code implementation of Model 2 with residual blocks and channel attention.

### 3.2.5 Our method: Model 3

Model 3 is an advanced Model with Perceptual Loss, Data Augmentation, and Learning Rate Scheduling

Building upon the previous advanced model, we further enhance the super resolution architecture by incorporating perceptual loss, data augmentation, and learning rate scheduling. These additions aim to improve the perceptual quality of the generated images, increase the model’s ro-

bustness to variations in the input data, and optimize the training process for better convergence and performance.

In addition to the mean absolute error (MAE) loss used in the previous models, Model 3 introduces perceptual loss [1] to capture high-level features and improve the visual quality of the super-resolved images. Perceptual loss is calculated using a pre-trained VGG19 network, which is known for its ability to extract meaningful features from images. The VGG19 network is used as a feature extractor, and the perceptual loss is computed as the mean squared error between the features of the ground truth and the generated images. By minimizing the perceptual loss, the model learns to generate images that are perceptually similar to the ground truth, resulting in sharper and more realistic details.

The choice to incorporate perceptual loss is motivated by the limitations of pixel-wise loss functions, such as MAE or mean squared error (MSE). These loss functions focus on minimizing the pixel-level differences between the generated and ground truth images, which can lead to blurry or overly smooth results. Perceptual loss, on the other hand, encourages the model to capture the high-level features and structures that are important for human perception. By combining perceptual loss with MAE loss, Model 3 strikes a balance between pixel-level accuracy and perceptual quality, resulting in super-resolved images that are both quantitatively and qualitatively superior.

To improve the model's ability to handle diverse input data and increase its robustness, Model 3 incorporates data augmentation techniques during training. Data augmentation involves applying random transformations to the training images, such as rotation, shifting, flipping, and scaling. By exposing the model to a wider range of variations in the input data, data augmentation helps the model learn more robust and generalizable features. In Model 3, we utilize the ImageDataGenerator from the Keras library to perform data augmentation. The ImageDataGenerator applies random rotations within a range of 20 degrees, random horizontal and vertical shifts of 10% of the image width and height, and random horizontal flips. These augmentations help the model learn to handle different orientations, positions, and reflections of the input images, making it more resilient to real-world variations. The decision to include data augmentation is based on the observation that the available training data may not cover all possible variations and transformations that can occur in real-world scenarios. By artificially expanding the training dataset through augmentation, we can improve the model's ability to generalize and handle a wider range of input conditions. Data augmentation also helps in reducing overfitting, as the model is exposed to a more diverse set of examples during training. Model 3 incorporates learning rate scheduling to optimize the training process and improve convergence. Learning rate scheduling involves adjusting the learning rate dynam-

cally during training based on a predefined schedule. In this model, we employ an exponential decay schedule, where the learning rate starts at an initial value and decays exponentially over time.

The exponential decay schedule is defined using the ExponentialDecay class from the Keras library. The initial learning rate is set to  $1e-4$ , and the decay rate is set to 0.95. The decay steps parameter determines the number of training steps after which the learning rate is decayed. By gradually reducing the learning rate, the model can converge more smoothly and avoid oscillations or divergence in the later stages of training.

The motivation behind using learning rate scheduling is to address the challenges associated with fixed learning rates. With a fixed learning rate, the model may take longer to converge or may get stuck in suboptimal solutions. By starting with a higher learning rate and gradually decaying it, the model can initially make larger steps in the parameter space, allowing it to explore and converge faster. As training progresses and the model approaches a good solution, the learning rate is reduced to enable finer adjustments and stabilize the convergence.

The architecture of Model 3 is similar to the previous advanced model, with the addition of more residual blocks. The model consists of a series of convolutional layers, followed by eight residual blocks, and then upsampling layers to generate the super-resolved image. The residual blocks help in capturing complex features and enabling deeper network training.

The model is trained using a combination of MAE loss and perceptual loss, with weights of 1.0 and 0.1, respectively. The Adam optimizer is used with the exponential decay learning rate schedule. The training is performed for 50 epochs, with a batch size of 4 and steps per epoch calculated based on the total number of training samples.

Model 3 is evaluated using the same metrics as the previous models, namely Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM). The model is tested on the validation set of the DIV2K dataset, and the average PSNR and SSIM scores are reported. Additionally, qualitative evaluation is performed by visually inspecting the super-resolved images to assess their perceptual quality, sharpness, and overall visual appeal.

By incorporating perceptual loss, data augmentation, and learning rate scheduling, Model 3 aims to generate super-resolved images that are both quantitatively and qualitatively superior to the previous models. The combination of these techniques helps in capturing perceptually relevant features, improving robustness to input variations, and optimizing the training process for better convergence and performance.

### 3.2.6 Final Model

The final model incorporates all the enhancements and techniques explored in the previous models, combining them into a comprehensive architecture for single image super resolution. This model builds upon the strengths of the prior approaches, aiming to generate high-quality, perceptually pleasing super-resolved images while addressing the limitations encountered earlier. The final model integrates channel attention mechanism [1] into the residual blocks to adaptively weigh the importance of different feature channels. Channel attention allows the model to focus on the most informative channels and suppress the less relevant ones, enhancing the model's ability to capture and reconstruct fine details. The channel attention module consists of a global average pooling layer followed by two fully connected layers with a bottleneck structure. The first fully connected layer reduces the channel dimension by a reduction ratio (set to 16 in this implementation), while the second layer restores the channel dimension. A sigmoid activation function is applied to generate channel-wise weights, which are then multiplied element-wise with the input feature maps. By incorporating channel attention, the final model can dynamically adjust the contribution of each channel based on its relevance to the super resolution task. This helps in prioritizing the most informative features and improving the model's ability to reconstruct high-frequency details. The final model employs a series of residual blocks [2] to facilitate the learning of deep representations and enable effective training of a deeper network. Each residual block consists of two convolutional layers with 3x3 filters, followed by LeakyReLU activation. A residual connection is added to allow the network to learn residual functions effectively. The model includes eight residual blocks, providing a deep architecture capable of capturing complex patterns and high-level features. The residual connections help in mitigating the vanishing gradient problem and allow for the efficient propagation of information across the network. The use of residual blocks in the final model enables the learning of more sophisticated representations, leading to improved super resolution performance compared to the previous models. In addition to the mean absolute error (MAE) loss, the final model incorporates perceptual loss [3] to enhance the perceptual quality of the generated images. Perceptual loss is computed using a pre-trained VGG19 network, which serves as a feature extractor. The perceptual loss is calculated as the mean squared error between the features extracted from the ground truth and the generated images using the VGG19 network. By minimizing the perceptual loss, the model learns to generate images that are perceptually similar to the ground truth, capturing high-level features and structures. The inclusion of perceptual loss helps in generating super-resolved images with sharper edges, finer textures, and improved visual quality. It com-

plements the pixel-wise MAE loss, striking a balance between pixel-level accuracy and perceptual fidelity.

To improve the model's robustness and generalization ability, the final model incorporates data augmentation techniques during training. Data augmentation involves applying random transformations to the training images, such as rotation, shifting, flipping, and scaling. The ImageDataGenerator from the Keras library is used to perform data augmentation. Random rotations within a range of 20 degrees, random horizontal and vertical shifts of 10% of the image width and height, and random horizontal flips are applied to the training images. Data augmentation helps in expanding the training dataset and exposing the model to a wider range of variations, enabling it to learn more robust and generalizable features. It reduces overfitting and improves the model's ability to handle diverse input conditions.

The final model employs learning rate scheduling to optimize the training process and improve convergence. An exponential decay schedule is used, where the learning rate starts at an initial value of  $1e-4$  and decays exponentially with a decay rate of 0.95 every 1000 steps. Learning rate scheduling allows the model to make larger steps in the parameter space during the initial stages of training, facilitating faster convergence. As training progresses, the learning rate is gradually reduced, enabling finer adjustments and stabilizing the convergence. The use of learning rate scheduling helps in finding a good balance between exploration and exploitation during the training process, leading to improved performance and stability.

The final model architecture consists of a series of convolutional layers, followed by eight residual blocks with channel attention, and then upsampling layers to generate the super-resolved image. The model takes a low-resolution image as input and progressively learns to reconstruct the corresponding high-resolution image. The model is trained using a combination of MAE loss and perceptual loss, with weights of 1.0 and 0.1, respectively. The Adam optimizer is used with the specified learning rate schedule. The training is performed for 100 epochs, with a batch size of 1 and steps per epoch calculated based on the total number of training samples.

The final model is evaluated using the DIV2K validation dataset. A custom evaluation function, `display_and_benchmark`, is implemented to assess the model's performance and compare it with the bicubic interpolation baseline. The evaluation function takes the validation dataset and the trained model as inputs. It iterates over a subset of the validation images (10 images in this case) and performs the following steps for each image:

- Upscales the low-resolution image using bicubic interpolation.

- Upscales the low-resolution image using the trained model.
- Calculates PSNR and SSIM metrics for both the bicubic interpolation and the model's output, comparing them with the ground truth high-resolution image.
- Displays a visual comparison of the low-resolution image, bicubic interpolation result, model's output, and the ground truth high-resolution image.
- Appends the PSNR and SSIM values to corresponding lists for later analysis.

After iterating over the subset of validation images, the average PSNR and SSIM values are calculated and printed for both the bicubic interpolation and the model's output. The `display_and_benchmark` function provides a comprehensive evaluation of the final model, assessing its performance in terms of quantitative metrics (PSNR and SSIM) and qualitative visual comparison. It allows for a direct comparison with the bicubic interpolation baseline, demonstrating the improvements achieved by the final model. By incorporating channel attention, residual blocks, perceptual loss, data augmentation, and learning rate scheduling, the final model aims to generate high-quality super-resolved images that exhibit sharp details, improved perceptual quality, and robustness to variations in the input data. The combination of these techniques enables the model to effectively learn deep representations, capture high-level features, and generate visually appealing results.

References: [1] Hu, J., Shen, L., Sun, G. (2018). Squeeze-and-excitation networks. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 7132-7141). [2] He, K., Zhang, X., Ren, S., Sun, J. (2016). Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 770-778). [3] Johnson, J., Alahi, A., Fei-Fei, L. (2016). Perceptual losses for real-time style transfer and super-resolution. In European conference on computer vision (pp. 694-711). Springer, Cham.

Appendix: See the provided code snippet for the implementation of the final model, including the channel attention module, residual blocks, perceptual loss, data augmentation, learning rate scheduling, and the evaluation function.



## 4. Results

### 4.1. Examples from Baseline

Below we compare examples of a low resolution image, an image upscaled through bicubic interpolation, and a high resolution image. Bicubic interpolation is an imperfect method for achieving super resolution, which can be clearly seen through the blurriness of edges and certain features.



### 4.2. Model 1: Initial Simple Deep Learning Results



Above are some of the results comparing a low resolution image, a solution using bicubic interpolation, initial our simple deep learning solution, and the high resolution image. While our simple deep learning solution sometimes did not perform as well as the bicubic solution on tests such as PSNR and SSIM, the comparisons seem to be fairly accurate to the naked eye. However, interestingly enough, our

Model	Avg. PSNR	Avg. SSIM
Master Model	35.44	0.926
Bicubic Interpolation	32.29	0.904
Model 3	34.00	0.927
Model 2	33.71	0.924
Model 1	4.80	0.025

Table 1. Average PSNR and SSIM values for different models.

solution seemed to have grid like patterns appear. This certainly effected the results for tests such as PSNR and SSIM, and were probably caused by the limitations of a simple DDPM solution.





Figure 1. Loss curves for different models over 50 epochs.

Quantitative results:



Solution	Average PSNR	Average SSIM
Bicubic	35.17	0.927
DDPM	35.44	0.93

Above are some of the results comparing a low resolution image, a solution using bicubic interpolation, our final simple deep learning solution, and the high resolution image. Here we were able to get our PNsR and SSIM scores up from adjusting parameters.






Figure 2. Loss curves for 100 epochs of training.

### 4.3. Model 3: Enhanced Deep Learning Model Results



Solution	Average PSNR	Average SSIM
Bicubic	34.97	0.917
Advanced DDPM	36.00	0.937

Above are some of the results comparing low resolution image, a solution using bicubic interpolation, our enhanced

deep learning solution using skip architecture, and the high resolution image. Using skip connections, we were able to score better on PSNR and SSIM, while using similar learning to our simple deep learning solution.

## 5. Conclusion

Ultimately, our novel project focuses on combining several aspects of Denoising Diffusion Probabilistic Models to other unique techniques we hypothesized may improve a deep learning model over bicubic interpolation. Namely, we use residual blocks in both the encoder and decoder to facilitate the learning of deep representations. We also utilize a shortcut connection to allow the network to learn residual functions effectively, which we hoped would mitigate any vanishing gradient problem and improve the flow of gradients during training. Additionally, we also incorporate skip connections between corresponding layers in the encoder and decoder to preserve spatial information and high-resolution details in the upsampling layers. We further employ a perceptual loss based on the VGG19 network in addition to the standard mean squared error (MSE) loss. Together, we hoped that these features would serve to capture more detail from low resolution images, leading to more detail in generated images.

However, since we did not use a pre-trained model, we found it difficult to reach an optimal balance between computational cost and training power. We managed to develop a model that ran relatively quickly on CPU, which was a choice made within the confines of a course project of this nature. If we had modified our code to run on GPU, access to which we would have needed to purchase, we predict that we could have dramatically improved our model performance through extensive training. Additionally, given more time, we would like to attempt to improve our PSNR score by training on augmented data. While we have architecture set up to do so, namely through data augmentation using flipping, cropping, and rotating, computational limitations of Google Colab prevented us from exploring how data augmentation would make our model more robust during training.

Despite these limitations, the loss curves for our models (Figures 1 and 2) demonstrate a general downward trend over the course of training, indicating that the models were learning effectively. Model 3, which incorporated perceptual loss, data augmentation, and learning rate scheduling, achieved the lowest loss values among our models, highlighting the benefits of these techniques.

Given that our advanced model was able to achieve a mild improvement over the bicubic interpolation baseline, in both PSNR and SSIM scores, we conclude that deep learning architectures are a meaningful avenue of further exploration in super resolution. Our basic deep learning model alone was able to achieve PSNR scores near the av-

erage for bicubic interpolation.

In conclusion, while we were unable to fully achieve our goal of generating high-resolution images with significantly enhanced detail and clarity compared to bicubic interpolation, our project serves as a valuable exploration of the potential of deep learning techniques for single image super resolution. The mild improvements achieved by our advanced models, along with the insights gained from our experiments, provide a foundation for future work in this area. Further refinement of the model architecture, extensive training on larger datasets, and the use of pre-trained models are potential avenues for improving upon our results and pushing the boundaries of deep learning-based super resolution.

## 6. References

Dong, C., & Loy, C. C. (2016). Image Super-Resolution Using Deep Convolutional Networks. *IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE*, VOL. 38(NO. 2).

Isola, P., Zhu, J. Y., Zhou, T., & Efros, A. A. (2017). Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 1125-1134).

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 770-778).

Hu, J., Shen, L., & Sun, G. (2018). Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 7132-7141).

Johnson, J., Alahi, A., & Fei-Fei, L. (2016). Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision* (pp. 694-711). Springer, Cham.

Krasser, M. (2018, October 17). Super-resolution/data.py at master · KRASSERM/Super-resolution. GitHub. <https://github.com/krasserm/super-resolution/blob/master/data.py>

Lim, B., Son, S., Kim, H., Nah, S., & Lee, K. M. (2017). Enhanced Deep Residual Networks for Single Image Super-Resolution. [https://doi.org/https://openaccess.thecvf.com/content\\_2017/w12/papers/Lim\\_2017.pdf](https://doi.org/https://openaccess.thecvf.com/content_2017/w12/papers/Lim_2017.pdf)

Meur, O. L., & Guillemot, C. (2012, January 1). Super-resolution-based inpainting. SpringerLink. [https://link.springer.com/chapter/10.1007/978-3-642-33783-3\\_40](https://link.springer.com/chapter/10.1007/978-3-642-33783-3_40)

Timofte, R., Agustsson, E., Gu, S., Wu, J., Ignatov, A., & Van Gool, L. (n.d.). Diverse 2K resolution high quality images as used for the challenges @ NTIRE (CVPR 2017 and CVPR 2018) and @ Pirm (ECCV 2018). DIV2K Dataset. <https://data.vision.ee.ethz.ch/cvl/DIV2K/>

## **7. Appendix**

See below



```
m tensorflow-addons) (24.0)
Collecting typeguard<3.0.0,>=2.7 (from tensorflow-addons)
  Downloading typeguard-2.13.3-py3-none-any.whl (17 kB)
Installing collected packages: typeguard, tensorflow-addons
Successfully installed tensorflow-addons-0.23.0 typeguard-2.13.3
```

```
In [ ]: import os
import sys
```

```
In [3]: import cv2
import numpy as np
import math
from skimage import metrics
from datasets import load_dataset
```

## BICUBIC INTERPOLATION IMPLEMENTATIONS

We have: CV2, GitHub and GitHub revised

```
In [4]: import cv2
def bicubic_interpolation_opencv(input_img, width, height):
    upscaled_img = cv2.resize(input_img, (width, height), interpolation=cv2.INTER_CUBIC)
    return upscaled_img
```

## DIV2K DATA LOADING AND OPERATIONS

We utilize publicly available code from the following github: <https://github.com/krasserm/super-resolution/blob/master/data.py>

```
In [5]: import os
import tensorflow as tf

from tensorflow.python.data.experimental import AUTOTUNE

class DIV2K:
    def __init__(self,
                 scale=2,
                 subset='train',
                 downgrade='bicubic',
                 images_dir='.div2k/images',
                 caches_dir='.div2k/caches'):

        self._ntire_2018 = True

        _scales = [2, 3, 4, 8]

        if scale in _scales:
            self.scale = scale
        else:
            raise ValueError(f'scale must be in ${_scales}')

        if subset == 'train':
            self.image_ids = range(1, 801)
        elif subset == 'valid':
            self.image_ids = range(801, 901)
        else:
            raise ValueError("subset must be 'train' or 'valid'")
```

```

_downgrades_a = ['bicubic', 'unknown']
_downgrades_b = ['mild', 'difficult']

if scale == 8 and downgrade != 'bicubic':
    raise ValueError(f'scale 8 only allowed for bicubic downgrade')

if downgrade in _downgrades_b and scale != 4:
    raise ValueError(f'{downgrade} downgrade requires scale 4')

if downgrade == 'bicubic' and scale == 8:
    self.downgrade = 'x8'
elif downgrade in _downgrades_b:
    self.downgrade = downgrade
else:
    self.downgrade = downgrade
    self._ntire_2018 = False

self.subset = subset
self.images_dir = images_dir
self.caches_dir = caches_dir

os.makedirs(images_dir, exist_ok=True)
os.makedirs(caches_dir, exist_ok=True)

def __len__(self):
    return len(self.image_ids)

def dataset(self, batch_size=16, repeat_count=None, random_transform=True):
    ds = tf.data.Dataset.zip((self.lr_dataset(), self.hr_dataset()))
    if random_transform:
        ds = ds.map(lambda lr, hr: random_crop(lr, hr, scale=self.scale), num_parallel_calls=AUTOTUNE)
        ds = ds.map(random_rotate, num_parallel_calls=AUTOTUNE)
        ds = ds.map(random_flip, num_parallel_calls=AUTOTUNE)
    ds = ds.batch(batch_size)
    ds = ds.repeat(repeat_count)
    ds = ds.prefetch(buffer_size=AUTOTUNE)
    return ds

def hr_dataset(self):
    if not os.path.exists(self._hr_images_dir()):
        download_archive(self._hr_images_archive(), self.images_dir, extract=True)

    ds = self._images_dataset(self._hr_image_files()).cache(self._hr_cache_file())

    if not os.path.exists(self._hr_cache_index()):
        self._populate_cache(ds, self._hr_cache_file())

    return ds

def lr_dataset(self):
    if not os.path.exists(self._lr_images_dir()):
        download_archive(self._lr_images_archive(), self.images_dir, extract=True)

    ds = self._images_dataset(self._lr_image_files()).cache(self._lr_cache_file())

    if not os.path.exists(self._lr_cache_index()):
        self._populate_cache(ds, self._lr_cache_file())

    return ds

def _hr_cache_file(self):
    return os.path.join(self.caches_dir, f'DIV2K_{self.subset}_HR.cache')

def _lr_cache_file(self):
    return os.path.join(self.caches_dir, f'DIV2K_{self.subset}_LR_{self.downgrade}_X

```

```

def _hr_cache_index(self):
    return f'{self._hr_cache_file()}.index'

def _lr_cache_index(self):
    return f'{self._lr_cache_file()}.index'

def _hr_image_files(self):
    images_dir = self._hr_images_dir()
    return [os.path.join(images_dir, f'{image_id:04}.png') for image_id in self.imag

def _lr_image_files(self):
    images_dir = self._lr_images_dir()
    return [os.path.join(images_dir, self._lr_image_file(image_id)) for image_id in

def _lr_image_file(self, image_id):
    if not self._ntire_2018 or self.scale == 8:
        return f'{image_id:04}x{self.scale}.png'
    else:
        return f'{image_id:04}x{self.scale}{self.downgrade[0]}.png'

def _hr_images_dir(self):
    return os.path.join(self.images_dir, f'DIV2K_{self.subset}_HR')

def _lr_images_dir(self):
    if self._ntire_2018:
        return os.path.join(self.images_dir, f'DIV2K_{self.subset}_LR_{self.downgrad
    else:
        return os.path.join(self.images_dir, f'DIV2K_{self.subset}_LR_{self.downgrad

def _hr_images_archive(self):
    return f'DIV2K_{self.subset}_HR.zip'

def _lr_images_archive(self):
    if self._ntire_2018:
        return f'DIV2K_{self.subset}_LR_{self.downgrade}.zip'
    else:
        return f'DIV2K_{self.subset}_LR_{self.downgrade}_X{self.scale}.zip'

@staticmethod
def _images_dataset(image_files):
    ds = tf.data.Dataset.from_tensor_slices(image_files)
    ds = ds.map(tf.io.read_file)
    ds = ds.map(lambda x: tf.image.decode_png(x, channels=3), num_parallel_calls=AUT
    return ds

@staticmethod
def _populate_cache(ds, cache_file):
    print(f'Caching decoded images in {cache_file} ...')
    for _ in ds: pass
    print(f'Cached decoded images in {cache_file}.')

# -----
# Transformations
# -----

def random_crop(lr_img, hr_img, hr_crop_size=96, scale=2):
    lr_crop_size = hr_crop_size // scale
    lr_img_shape = tf.shape(lr_img)[:2]

    lr_w = tf.random.uniform(shape=(), maxval=lr_img_shape[1] - lr_crop_size + 1, dtype=
    lr_h = tf.random.uniform(shape=(), maxval=lr_img_shape[0] - lr_crop_size + 1, dtype=

    hr_w = lr_w * scale

```

```

hr_h = lr_h * scale

lr_img_cropped = lr_img[lr_h:lr_h + lr_crop_size, lr_w:lr_w + lr_crop_size]
hr_img_cropped = hr_img[hr_h:hr_h + hr_crop_size, hr_w:hr_w + hr_crop_size]

return lr_img_cropped, hr_img_cropped

def random_flip(lr_img, hr_img):
    rn = tf.random.uniform(shape=(), maxval=1)
    return tf.cond(rn < 0.5,
                   lambda: (lr_img, hr_img),
                   lambda: (tf.image.flip_left_right(lr_img),
                             tf.image.flip_left_right(hr_img)))

def random_rotate(lr_img, hr_img):
    rn = tf.random.uniform(shape=(), maxval=4, dtype=tf.int32)
    return tf.image.rot90(lr_img, rn), tf.image.rot90(hr_img, rn)

# -----
# IO
# -----

def download_archive(file, target_dir, extract=True):
    source_url = f'http://data.vision.ee.ethz.ch/cvl/DIV2K/{file}'
    target_dir = os.path.abspath(target_dir)
    tf.keras.utils.get_file(file, source_url, cache_subdir=target_dir, extract=extract)
    os.remove(os.path.join(target_dir, file))

```

## MODEL 1 Base

```

In [12]: import tensorflow as tf
from tensorflow.keras.layers import Input, Conv2D, Conv2DTranspose, LeakyReLU, Lambda, A
from tensorflow.keras.models import Model
from tensorflow.keras.optimizers import Adam
import numpy as np
from skimage.metrics import structural_similarity as ssim
from skimage.metrics import peak_signal_noise_ratio as psnr

def channel_attention(x, reduction_ratio=16):
    _, _, _, c = x.shape
    avg_pool = GlobalAveragePooling2D()(x)
    avg_pool = Dense(c // reduction_ratio, activation='relu')(avg_pool)
    avg_pool = Dense(c, activation='sigmoid')(avg_pool)
    return Multiply()([x, avg_pool])

def residual_block(x, filters):
    residual = x
    x = Conv2D(filters, (3, 3), padding='same')(x)
    x = LeakyReLU(alpha=0.2)(x)
    x = Conv2D(filters, (3, 3), padding='same')(x)
    x = channel_attention(x)
    x = Add()([x, residual])
    x = LeakyReLU(alpha=0.2)(x)
    return x

def create_model(scale=2):
    inputs = tf.keras.Input(shape=(None, None, 3))
    x = Conv2D(64, (5, 5), padding='same')(inputs)
    x = LeakyReLU(alpha=0.2)(x)

```

```

for _ in range(6):
    x = residual_block(x, 64)

x = Conv2D(64 * (scale ** 2), (3, 3), padding='same')(x)
x = tf.nn.depth_to_space(x, scale)
x = LeakyReLU(alpha=0.2)(x)
x = Conv2D(3, (5, 5), padding='same')(x)

outputs = Add()(x, tf.keras.layers.UpSampling2D(size=(scale, scale), interpolation=
model = Model(inputs, outputs)
return model

scale = 2
model = create_model(scale)
model.compile(optimizer=Adam(learning_rate=1e-4), loss='mean_absolute_error')

val_div2k = DIV2K(scale=scale, subset='valid', downgrade='bicubic')
val_dataset = val_div2k.dataset(batch_size=1, repeat_count=1)
val_dataset = val_dataset.prefetch(tf.data.experimental.AUTOTUNE)

total_samples = len(val_div2k)
batch_size = 1
steps_per_epoch = total_samples // batch_size
model.fit(val_dataset, epochs=100, steps_per_epoch=steps_per_epoch)

psnr_values = []
ssim_values = []
for lr, hr in val_dataset:
    sr = model.predict(lr)
    hr = hr.numpy().squeeze()
    sr = sr.squeeze()

    psnr_value = psnr(hr, sr)
    ssim_value = ssim(hr, sr, multichannel=True)

    psnr_values.append(psnr_value)
    ssim_values.append(ssim_value)

print(f"Average PSNR: {np.mean(psnr_values)}")
print(f"Average SSIM: {np.mean(ssim_values)}")

```

```

Epoch 1/100
100/100 [=====] - 10s 18ms/step - loss: 6.6328
Epoch 2/100
100/100 [=====] - 2s 16ms/step - loss: 5.2139
Epoch 3/100
100/100 [=====] - 2s 16ms/step - loss: 4.4872
Epoch 4/100
100/100 [=====] - 2s 16ms/step - loss: 5.1317
Epoch 5/100
100/100 [=====] - 2s 16ms/step - loss: 4.4164
Epoch 6/100
100/100 [=====] - 2s 24ms/step - loss: 4.6974
Epoch 7/100
100/100 [=====] - 2s 18ms/step - loss: 5.0067
Epoch 8/100
100/100 [=====] - 2s 18ms/step - loss: 4.6284
Epoch 9/100
100/100 [=====] - 2s 17ms/step - loss: 4.3068
Epoch 10/100
100/100 [=====] - 2s 16ms/step - loss: 5.1852
Epoch 11/100
100/100 [=====] - 2s 22ms/step - loss: 4.3788
Epoch 12/100
100/100 [=====] - 2s 16ms/step - loss: 4.1929

```

Epoch 13/100  
100/100 [=====] - 2s 16ms/step - loss: 4.9578  
Epoch 14/100  
100/100 [=====] - 2s 16ms/step - loss: 4.6085  
Epoch 15/100  
100/100 [=====] - 2s 16ms/step - loss: 4.2786  
Epoch 16/100  
100/100 [=====] - 2s 16ms/step - loss: 4.6594  
Epoch 17/100  
100/100 [=====] - 2s 20ms/step - loss: 4.3764  
Epoch 18/100  
100/100 [=====] - 2s 16ms/step - loss: 4.1494  
Epoch 19/100  
100/100 [=====] - 2s 17ms/step - loss: 4.7503  
Epoch 20/100  
100/100 [=====] - 2s 18ms/step - loss: 4.5712  
Epoch 21/100  
100/100 [=====] - 2s 18ms/step - loss: 4.5654  
Epoch 22/100  
100/100 [=====] - 2s 20ms/step - loss: 4.6715  
Epoch 23/100  
100/100 [=====] - 2s 17ms/step - loss: 4.8929  
Epoch 24/100  
100/100 [=====] - 2s 16ms/step - loss: 4.1877  
Epoch 25/100  
100/100 [=====] - 2s 16ms/step - loss: 4.2544  
Epoch 26/100  
100/100 [=====] - 2s 17ms/step - loss: 4.1167  
Epoch 27/100  
100/100 [=====] - 2s 22ms/step - loss: 4.5573  
Epoch 28/100  
100/100 [=====] - 2s 16ms/step - loss: 4.0833  
Epoch 29/100  
100/100 [=====] - 2s 17ms/step - loss: 4.4788  
Epoch 30/100  
100/100 [=====] - 2s 16ms/step - loss: 4.6313  
Epoch 31/100  
100/100 [=====] - 2s 16ms/step - loss: 4.6510  
Epoch 32/100  
100/100 [=====] - 2s 16ms/step - loss: 4.2262  
Epoch 33/100  
100/100 [=====] - 2s 20ms/step - loss: 4.7140  
Epoch 34/100  
100/100 [=====] - 2s 18ms/step - loss: 4.4767  
Epoch 35/100  
100/100 [=====] - 2s 19ms/step - loss: 4.4583  
Epoch 36/100  
100/100 [=====] - 2s 17ms/step - loss: 4.2397  
Epoch 37/100  
100/100 [=====] - 2s 16ms/step - loss: 4.5085  
Epoch 38/100  
100/100 [=====] - 2s 17ms/step - loss: 4.2894  
Epoch 39/100  
100/100 [=====] - 2s 21ms/step - loss: 4.0633  
Epoch 40/100  
100/100 [=====] - 2s 17ms/step - loss: 4.5077  
Epoch 41/100  
100/100 [=====] - 2s 16ms/step - loss: 4.0260  
Epoch 42/100  
100/100 [=====] - 2s 16ms/step - loss: 4.8251  
Epoch 43/100  
100/100 [=====] - 2s 16ms/step - loss: 4.2361  
Epoch 44/100  
100/100 [=====] - 2s 24ms/step - loss: 4.1534  
Epoch 45/100  
100/100 [=====] - 2s 17ms/step - loss: 3.7210

Epoch 46/100  
100/100 [=====] - 2s 16ms/step - loss: 3.8551  
Epoch 47/100  
100/100 [=====] - 2s 16ms/step - loss: 4.3404  
Epoch 48/100  
100/100 [=====] - 2s 16ms/step - loss: 4.1247  
Epoch 49/100  
100/100 [=====] - 2s 17ms/step - loss: 3.9264  
Epoch 50/100  
100/100 [=====] - 2s 19ms/step - loss: 3.8333  
Epoch 51/100  
100/100 [=====] - 2s 17ms/step - loss: 4.2724  
Epoch 52/100  
100/100 [=====] - 2s 16ms/step - loss: 4.1318  
Epoch 53/100  
100/100 [=====] - 2s 16ms/step - loss: 3.6507  
Epoch 54/100  
100/100 [=====] - 2s 16ms/step - loss: 3.8927  
Epoch 55/100  
100/100 [=====] - 2s 19ms/step - loss: 3.6430  
Epoch 56/100  
100/100 [=====] - 2s 16ms/step - loss: 3.8063  
Epoch 57/100  
100/100 [=====] - 2s 16ms/step - loss: 4.2884  
Epoch 58/100  
100/100 [=====] - 2s 16ms/step - loss: 4.1293  
Epoch 59/100  
100/100 [=====] - 2s 16ms/step - loss: 3.6993  
Epoch 60/100  
100/100 [=====] - 2s 17ms/step - loss: 3.5527  
Epoch 61/100  
100/100 [=====] - 2s 22ms/step - loss: 4.0577  
Epoch 62/100  
100/100 [=====] - 2s 17ms/step - loss: 3.8931  
Epoch 63/100  
100/100 [=====] - 2s 16ms/step - loss: 3.9954  
Epoch 64/100  
100/100 [=====] - 2s 16ms/step - loss: 4.2498  
Epoch 65/100  
100/100 [=====] - 2s 16ms/step - loss: 3.8123  
Epoch 66/100  
100/100 [=====] - 2s 22ms/step - loss: 4.0389  
Epoch 67/100  
100/100 [=====] - 2s 16ms/step - loss: 4.1420  
Epoch 68/100  
100/100 [=====] - 2s 16ms/step - loss: 4.1216  
Epoch 69/100  
100/100 [=====] - 2s 18ms/step - loss: 4.4848  
Epoch 70/100  
100/100 [=====] - 2s 24ms/step - loss: 3.5785  
Epoch 71/100  
100/100 [=====] - 2s 18ms/step - loss: 4.2249  
Epoch 72/100  
100/100 [=====] - 2s 21ms/step - loss: 4.4478  
Epoch 73/100  
100/100 [=====] - 2s 17ms/step - loss: 4.3258  
Epoch 74/100  
100/100 [=====] - 2s 16ms/step - loss: 4.1413  
Epoch 75/100  
100/100 [=====] - 2s 16ms/step - loss: 3.7085  
Epoch 76/100  
100/100 [=====] - 2s 16ms/step - loss: 3.8050  
Epoch 77/100  
100/100 [=====] - 2s 21ms/step - loss: 4.2007  
Epoch 78/100  
100/100 [=====] - 2s 19ms/step - loss: 4.1449

```
Epoch 79/100
100/100 [=====] - 2s 16ms/step - loss: 4.1545
Epoch 80/100
100/100 [=====] - 2s 16ms/step - loss: 4.0690
Epoch 81/100
100/100 [=====] - 2s 16ms/step - loss: 3.7793
Epoch 82/100
100/100 [=====] - 2s 16ms/step - loss: 3.8082
Epoch 83/100
100/100 [=====] - 2s 24ms/step - loss: 3.6327
Epoch 84/100
100/100 [=====] - 2s 16ms/step - loss: 4.0491
Epoch 85/100
100/100 [=====] - 2s 16ms/step - loss: 3.9945
Epoch 86/100
100/100 [=====] - 2s 17ms/step - loss: 3.9046
Epoch 87/100
100/100 [=====] - 2s 16ms/step - loss: 3.6986
Epoch 88/100
100/100 [=====] - 2s 17ms/step - loss: 3.6224
Epoch 89/100
100/100 [=====] - 2s 21ms/step - loss: 3.5070
Epoch 90/100
100/100 [=====] - 2s 17ms/step - loss: 4.2493
Epoch 91/100
100/100 [=====] - 2s 16ms/step - loss: 3.7345
Epoch 92/100
100/100 [=====] - 2s 16ms/step - loss: 3.7521
Epoch 93/100
100/100 [=====] - 2s 17ms/step - loss: 4.3850
Epoch 94/100
100/100 [=====] - 2s 18ms/step - loss: 3.8579
Epoch 95/100
100/100 [=====] - 2s 21ms/step - loss: 3.8502
Epoch 96/100
100/100 [=====] - 2s 16ms/step - loss: 3.8819
Epoch 97/100
100/100 [=====] - 2s 16ms/step - loss: 3.8489
Epoch 98/100
100/100 [=====] - 2s 15ms/step - loss: 3.9357
Epoch 99/100
100/100 [=====] - 2s 16ms/step - loss: 3.8882
Epoch 100/100
100/100 [=====] - 2s 16ms/step - loss: 4.2511
1/1 [=====] - 0s 451ms/step
1/1 [=====] - 0s 32ms/step
1/1 [=====] - 0s 29ms/step
```

```
<ipython-input-12-b49ba6614d53>:63: UserWarning: Inputs have mismatched dtype. Setting data_range based on image_true.
```

```
    psnr_value = psnr(hr, sr)
```

```
<ipython-input-12-b49ba6614d53>:64: FutureWarning: `multichannel` is a deprecated argument name for `structural_similarity`. It will be removed in version 1.0. Please use `channel_axis` instead.
```

```
    ssim_value = ssim(hr, sr, multichannel=True)
```

```
/usr/local/lib/python3.10/dist-packages/skimage/_shared/utils.py:348: UserWarning: Inputs have mismatched dtype. Setting data_range based on im1.dtype.
```

```
    return func(*args, **kwargs)
```

```
1/1 [=====] - 0s 28ms/step
1/1 [=====] - 0s 28ms/step
1/1 [=====] - 0s 28ms/step
1/1 [=====] - 0s 27ms/step
1/1 [=====] - 0s 30ms/step
1/1 [=====] - 0s 33ms/step
1/1 [=====] - 0s 33ms/step
1/1 [=====] - 0s 35ms/step
```



```

1/1 [=====] - 0s 18ms/step
1/1 [=====] - 0s 18ms/step
1/1 [=====] - 0s 20ms/step
1/1 [=====] - 0s 23ms/step
1/1 [=====] - 0s 18ms/step
1/1 [=====] - 0s 18ms/step
1/1 [=====] - 0s 22ms/step
1/1 [=====] - 0s 20ms/step
1/1 [=====] - 0s 20ms/step
1/1 [=====] - 0s 19ms/step
1/1 [=====] - 0s 19ms/step
1/1 [=====] - 0s 21ms/step
1/1 [=====] - 0s 21ms/step
1/1 [=====] - 0s 18ms/step
1/1 [=====] - 0s 20ms/step
1/1 [=====] - 0s 18ms/step

```

Average PSNR: 35.43503249328994  
Average SSIM: 0.9259011426940108

```

In [ ]: import tensorflow as tf
from tensorflow.keras.layers import Conv2D, LeakyReLU, Add, UpSampling2D
from tensorflow.keras.models import Model
from tensorflow.keras.applications import VGG19
from tensorflow.keras.optimizers import Adam
from tensorflow.keras.preprocessing.image import ImageDataGenerator
from tensorflow.keras.optimizers.schedules import ExponentialDecay

from skimage.metrics import structural_similarity as ssim
from skimage.metrics import peak_signal_noise_ratio as psnr

def residual_block(x, filters):
    residual = x
    x = Conv2D(filters, (3, 3), padding='same')(x)
    x = LeakyReLU(alpha=0.2)(x)
    x = Conv2D(filters, (3, 3), padding='same')(x)
    x = Add()([x, residual])
    x = LeakyReLU(alpha=0.2)(x)
    return x

def create_model(scale=2):
    inputs = tf.keras.Input(shape=(None, None, 3))
    x = Conv2D(128, (5, 5), padding='same')(inputs)
    x = LeakyReLU(alpha=0.2)(x)

    for _ in range(8):
        x = residual_block(x, 128)

    x = Conv2D(128 * (scale ** 2), (3, 3), padding='same')(x)
    x = tf.nn.depth_to_space(x, scale)
    x = LeakyReLU(alpha=0.2)(x)
    x = Conv2D(3, (5, 5), padding='same')(x)

    outputs = Add()([x, tf.keras.layers.UpSampling2D(size=(scale, scale), interpolation=
    model = Model(inputs, outputs)
    return model

vgg = VGG19(weights='imagenet', include_top=False, input_shape=(None, None, 3))
vgg.trainable = False

```

```

def perceptual_loss(y_true, y_pred):
    vgg_true = vgg(y_true)
    vgg_pred = vgg(y_pred)
    return tf.reduce_mean(tf.square(vgg_true - vgg_pred))

model.compile(optimizer=Adam(learning_rate=1e-4), loss=['mean_absolute_error', perceptua

batch_size = 4
steps_per_epoch = total_samples // batch_size

datagen = ImageDataGenerator(
    rotation_range=20,
    width_shift_range=0.1,
    height_shift_range=0.1,
    horizontal_flip=True
)

lr_schedule = tf.keras.optimizers.schedules.ExponentialDecay(
    initial_learning_rate=1e-4,
    decay_steps=1000,
    decay_rate=0.95,
    staircase=True
)
optimizer = Adam(learning_rate=lr_schedule)

scale = 2
model = create_model(scale)
model.compile(optimizer=Adam(learning_rate=1e-4), loss='mean_absolute_error')

val_div2k = DIV2K(scale=scale, subset='valid', downgrade='bicubic')
val_dataset = val_div2k.dataset(batch_size=1, repeat_count=1)
val_dataset = val_dataset.prefetch(tf.data.experimental.AUTOTUNE)

total_samples = len(val_div2k)
batch_size = 1
steps_per_epoch = total_samples // batch_size
model.fit(val_dataset, epochs=50, steps_per_epoch=steps_per_epoch)

psnr_values = []
ssim_values = []
for lr, hr in val_dataset:
    sr = model.predict(lr)
    hr = hr.numpy().squeeze()
    sr = sr.squeeze()

    psnr_value = psnr(hr, sr)
    ssim_value = ssim(hr, sr, multichannel=True)

    psnr_values.append(psnr_value)
    ssim_values.append(ssim_value)

print(f"Average PSNR: {np.mean(psnr_values)}")
print(f"Average SSIM: {np.mean(ssim_values)}")

```

## SKIP CONNECTIONS ON DDPM

## NEW SKIP ARCHITECTURE

### 1. Model Definition

Input Layer: Takes an image of any size with three channels (RGB).

Encoder: Consists of several convolutional layers (using Conv2D) with increasing number of filters and strides to downsample the image, processing it into deeper feature representations. Each layer is followed by a LeakyReLU activation for non-linearity and several custom residual blocks.

Decoder: Uses Conv2DTranspose for upsampling, doubling the resolution with each step. Concatenation with corresponding encoder feature maps (Concatenate()) suggests a U-Net-like architecture, which helps in recovering fine details by combining low-level and high-level features.

<https://medium.com/@danushidk507/skip-connections-ab515d634e6d>

Output Layer: Produces the final image with the same resolution as the input but presumably enhanced, as indicated by the tanh activation, which normalizes the output pixel values between -1 and 1.

1. Loss Function Perceptual Loss: A secondary loss function defined using the VGG19 model, focusing on the similarity of deep features between the predicted and true images, which is beneficial for maintaining textural details.
1. Training and Evaluation Data Preparation: Uses a DIV2K dataset specifically prepared for a super-resolution task with a scaling factor of 2, indicating that the target is to double the resolution of input images.

Training Procedure: Model is trained using a mixed loss function that combines Mean Squared Error (MSE) for pixel-wise accuracy and perceptual loss for maintaining textural details.

Learning Rate Scheduler: Employs an exponential decay schedule for adjusting the learning rate, which helps in stabilizing the training as it progresses. Metrics Calculation: After training, the model evaluates the performance using PSNR (Peak Signal-to-Noise Ratio) and SSIM (Structural Similarity Index Measure), both are standard metrics for assessing image quality in tasks like super-resolution.

```
In [20]: import tensorflow as tf
from tensorflow.keras.layers import Conv2D, Conv2DTranspose, LeakyReLU, Add, Lambda, Mul
from tensorflow.keras.models import Model
from tensorflow.keras.applications import VGG19
import matplotlib.pyplot as plt
from skimage.metrics import peak_signal_noise_ratio as psnr
from skimage.metrics import structural_similarity as ssim

def channel_attention(x, reduction_ratio=16):
    _, _, _, c = x.shape
    avg_pool = GlobalAveragePooling2D()(x)
    avg_pool = Dense(c // reduction_ratio, activation='relu')(avg_pool)
    avg_pool = Dense(c, activation='sigmoid')(avg_pool)
    return Multiply()([x, avg_pool])

def residual_block(x, filters):
    residual = x
    x = Conv2D(filters, (3, 3), padding='same')(x)
    x = LeakyReLU(alpha=0.2)(x)
    x = Conv2D(filters, (3, 3), padding='same')(x)
    x = channel_attention(x)
    x = Add()([x, residual])
    x = LeakyReLU(alpha=0.2)(x)
    return x

def create_model(scale=2):
    inputs = tf.keras.Input(shape=(None, None, 3))
```

```

x = Conv2D(128, (5, 5), padding='same')(inputs)
x = LeakyReLU(alpha=0.2)(x)

for _ in range(8):
    x = residual_block(x, 128)

x = Conv2D(128 * (scale ** 2), (3, 3), padding='same')(x)
x = tf.nn.depth_to_space(x, scale)
x = LeakyReLU(alpha=0.2)(x)
x = Conv2D(3, (5, 5), padding='same')(x)

outputs = Add()(x, tf.keras.layers.UpSampling2D(size=(scale, scale), interpolation=
model = Model(inputs, outputs)
return model

vgg = VGG19(weights='imagenet', include_top=False, input_shape=(None, None, 3))
vgg.trainable = False

def perceptual_loss(y_true, y_pred):
    vgg_true = vgg(y_true)
    vgg_pred = vgg(y_pred)
    return tf.reduce_mean(tf.square(vgg_true - vgg_pred))

model.compile(optimizer=Adam(learning_rate=1e-4), loss=['mean_absolute_error', perceptua

batch_size = 4
steps_per_epoch = total_samples // batch_size

datagen = ImageDataGenerator(
    rotation_range=20,
    width_shift_range=0.1,
    height_shift_range=0.1,
    horizontal_flip=True
)

lr_schedule = tf.keras.optimizers.schedules.ExponentialDecay(
    initial_learning_rate=1e-4,
    decay_steps=1000,
    decay_rate=0.95,
    staircase=True
)
optimizer = Adam(learning_rate=lr_schedule)

scale = 2
model = create_model(scale)
model.compile(optimizer=tf.keras.optimizers.Adam(learning_rate=1e-4), loss=['mean_absolu

val_div2k = DIV2K(scale=scale, subset='valid', downgrade='bicubic')
val_dataset = val_div2k.dataset(batch_size=1, repeat_count=1)
val_dataset = val_dataset.prefetch(tf.data.experimental.AUTOTUNE)

total_samples = len(val_div2k)
batch_size = 1
steps_per_epoch = total_samples // batch_size

model.fit(val_dataset, epochs=100, steps_per_epoch=steps_per_epoch)

def display_and_benchmark(dataset, model):
    psnr_values_bicubic = []
    ssim_values_bicubic = []
    psnr_values_model = []
    ssim_values_model = []

    for lr_img, hr_img in dataset.take(10):
        upscaled_img_bicubic = tf.image.resize(lr_img, [hr_img.shape[1], hr_img.shape[2]
            if upscaled_img_bicubic.dtype != tf.uint8:

```

```

upscaled_img_bicubic = tf.clip_by_value(upscaled_img_bicubic, 0.0, 255.0)
upscaled_img_bicubic = tf.cast(upscaled_img_bicubic, tf.uint8)

upscaled_img_model = model.predict(lr_img)
if upscaled_img_model.dtype != tf.uint8:
    upscaled_img_model = tf.clip_by_value(upscaled_img_model, 0.0, 255.0)
    upscaled_img_model = tf.cast(upscaled_img_model, tf.uint8)

hr_img_np = hr_img.numpy()[0]
upscaled_img_bicubic_np = upscaled_img_bicubic.numpy()[0]
upscaled_img_model_np = upscaled_img_model.numpy()[0]

current_psnr_bicubic = psnr(hr_img_np, upscaled_img_bicubic_np, data_range=hr_img_np)
current_ssim_bicubic = ssim(hr_img_np, upscaled_img_bicubic_np, multichannel=True)
psnr_values_bicubic.append(current_psnr_bicubic)
ssim_values_bicubic.append(current_ssim_bicubic)

current_psnr_model = psnr(hr_img_np, upscaled_img_model_np, data_range=hr_img_np)
current_ssim_model = ssim(hr_img_np, upscaled_img_model_np, multichannel=True)
psnr_values_model.append(current_psnr_model)
ssim_values_model.append(current_ssim_model)

plt.figure(figsize=(18, 6))
plt.subplot(1, 4, 1)
plt.imshow(lr_img.numpy()[0])
plt.title('Low-Resolution')
plt.axis('off')

plt.subplot(1, 4, 2)
plt.imshow(upscaled_img_bicubic_np)
plt.title(f'Bicubic\nPSNR: {current_psnr_bicubic:.2f}, SSIM: {current_ssim_bicubic:.3f}')
plt.axis('off')

plt.subplot(1, 4, 3)
plt.imshow(upscaled_img_model_np)
plt.title(f'Model\nPSNR: {current_psnr_model:.2f}, SSIM: {current_ssim_model:.3f}')
plt.axis('off')

plt.subplot(1, 4, 4)
plt.imshow(hr_img_np)
plt.title('High-Resolution')
plt.axis('off')

plt.tight_layout()
plt.show()

print(f'Average PSNR (Bicubic): {np.mean(psnr_values_bicubic):.2f}')
print(f'Average SSIM (Bicubic): {np.mean(ssim_values_bicubic):.3f}')
print(f'Average PSNR (Model): {np.mean(psnr_values_model):.2f}')
print(f'Average SSIM (Model): {np.mean(ssim_values_model):.3f}')

display_and_benchmark(val_dataset, model)

```

```

Epoch 1/100
100/100 [=====] - 12s 21ms/step - loss: 6.1537
Epoch 2/100
100/100 [=====] - 3s 26ms/step - loss: 5.2918
Epoch 3/100
100/100 [=====] - 2s 20ms/step - loss: 4.8333
Epoch 4/100
100/100 [=====] - 2s 20ms/step - loss: 4.4832
Epoch 5/100
100/100 [=====] - 2s 20ms/step - loss: 5.1404
Epoch 6/100
100/100 [=====] - 2s 20ms/step - loss: 4.4439
Epoch 7/100

```

```
100/100 [=====] - 3s 26ms/step - loss: 5.0036
Epoch 8/100
100/100 [=====] - 2s 20ms/step - loss: 4.4763
Epoch 9/100
100/100 [=====] - 2s 20ms/step - loss: 4.9316
Epoch 10/100
100/100 [=====] - 2s 21ms/step - loss: 4.2035
Epoch 11/100
100/100 [=====] - 2s 23ms/step - loss: 4.3652
Epoch 12/100
100/100 [=====] - 2s 24ms/step - loss: 3.9943
Epoch 13/100
100/100 [=====] - 2s 20ms/step - loss: 4.4427
Epoch 14/100
100/100 [=====] - 2s 20ms/step - loss: 4.2816
Epoch 15/100
100/100 [=====] - 2s 21ms/step - loss: 3.5615
Epoch 16/100
100/100 [=====] - 2s 24ms/step - loss: 4.2047
Epoch 17/100
100/100 [=====] - 2s 23ms/step - loss: 4.5697
Epoch 18/100
100/100 [=====] - 2s 20ms/step - loss: 4.1477
Epoch 19/100
100/100 [=====] - 2s 21ms/step - loss: 3.8330
Epoch 20/100
100/100 [=====] - 2s 20ms/step - loss: 4.2806
Epoch 21/100
100/100 [=====] - 3s 30ms/step - loss: 4.0464
Epoch 22/100
100/100 [=====] - 2s 21ms/step - loss: 4.0636
Epoch 23/100
100/100 [=====] - 2s 21ms/step - loss: 4.3112
Epoch 24/100
100/100 [=====] - 2s 21ms/step - loss: 4.2138
Epoch 25/100
100/100 [=====] - 2s 24ms/step - loss: 4.3040
Epoch 26/100
100/100 [=====] - 2s 24ms/step - loss: 3.9627
Epoch 27/100
100/100 [=====] - 2s 20ms/step - loss: 3.4976
Epoch 28/100
100/100 [=====] - 2s 21ms/step - loss: 4.3168
Epoch 29/100
100/100 [=====] - 2s 20ms/step - loss: 4.2669
Epoch 30/100
100/100 [=====] - 2s 23ms/step - loss: 4.0614
Epoch 31/100
100/100 [=====] - 2s 23ms/step - loss: 4.0492
Epoch 32/100
100/100 [=====] - 2s 20ms/step - loss: 3.8667
Epoch 33/100
100/100 [=====] - 2s 21ms/step - loss: 4.7078
Epoch 34/100
100/100 [=====] - 2s 20ms/step - loss: 3.8092
Epoch 35/100
100/100 [=====] - 3s 31ms/step - loss: 3.8190
Epoch 36/100
100/100 [=====] - 2s 23ms/step - loss: 3.7577
Epoch 37/100
100/100 [=====] - 2s 21ms/step - loss: 3.9243
Epoch 38/100
100/100 [=====] - 2s 20ms/step - loss: 4.4415
Epoch 39/100
100/100 [=====] - 2s 20ms/step - loss: 4.0058
Epoch 40/100
```

```
100/100 [=====] - 2s 22ms/step - loss: 4.2212
Epoch 41/100
100/100 [=====] - 2s 23ms/step - loss: 3.8746
Epoch 42/100
100/100 [=====] - 2s 20ms/step - loss: 3.5312
Epoch 43/100
100/100 [=====] - 2s 20ms/step - loss: 4.0531
Epoch 44/100
100/100 [=====] - 2s 21ms/step - loss: 3.9304
Epoch 45/100
100/100 [=====] - 3s 30ms/step - loss: 4.0197
Epoch 46/100
100/100 [=====] - 2s 21ms/step - loss: 4.0842
Epoch 47/100
100/100 [=====] - 2s 20ms/step - loss: 3.7493
Epoch 48/100
100/100 [=====] - 3s 33ms/step - loss: 3.8759
Epoch 49/100
100/100 [=====] - 2s 21ms/step - loss: 4.4397
Epoch 50/100
100/100 [=====] - 2s 21ms/step - loss: 3.7914
Epoch 51/100
100/100 [=====] - 2s 20ms/step - loss: 3.5741
Epoch 52/100
100/100 [=====] - 2s 20ms/step - loss: 4.1056
Epoch 53/100
100/100 [=====] - 3s 27ms/step - loss: 4.0586
Epoch 54/100
100/100 [=====] - 2s 21ms/step - loss: 3.7243
Epoch 55/100
100/100 [=====] - 2s 20ms/step - loss: 3.9936
Epoch 56/100
100/100 [=====] - 2s 20ms/step - loss: 3.5135
Epoch 57/100
100/100 [=====] - 2s 25ms/step - loss: 3.9543
Epoch 58/100
100/100 [=====] - 2s 20ms/step - loss: 4.2004
Epoch 59/100
100/100 [=====] - 2s 20ms/step - loss: 4.4225
Epoch 60/100
100/100 [=====] - 2s 20ms/step - loss: 4.4181
Epoch 61/100
100/100 [=====] - 2s 20ms/step - loss: 4.0670
Epoch 62/100
100/100 [=====] - 3s 27ms/step - loss: 4.3106
Epoch 63/100
100/100 [=====] - 2s 20ms/step - loss: 3.7085
Epoch 64/100
100/100 [=====] - 2s 20ms/step - loss: 3.7898
Epoch 65/100
100/100 [=====] - 2s 20ms/step - loss: 3.9510
Epoch 66/100
100/100 [=====] - 2s 20ms/step - loss: 3.7653
Epoch 67/100
100/100 [=====] - 3s 27ms/step - loss: 3.6184
Epoch 68/100
100/100 [=====] - 2s 20ms/step - loss: 4.2421
Epoch 69/100
100/100 [=====] - 2s 21ms/step - loss: 3.8267
Epoch 70/100
100/100 [=====] - 2s 20ms/step - loss: 3.8067
Epoch 71/100
100/100 [=====] - 2s 20ms/step - loss: 3.7800
Epoch 72/100
100/100 [=====] - 3s 27ms/step - loss: 3.9919
Epoch 73/100
```


```
100/100 [=====] - 2s 20ms/step - loss: 4.3193
Epoch 74/100
100/100 [=====] - 2s 20ms/step - loss: 3.8108
Epoch 75/100
100/100 [=====] - 2s 20ms/step - loss: 3.8260
Epoch 76/100
100/100 [=====] - 3s 27ms/step - loss: 3.8263
Epoch 77/100
100/100 [=====] - 2s 21ms/step - loss: 3.7684
Epoch 78/100
100/100 [=====] - 2s 20ms/step - loss: 4.0402
Epoch 79/100
100/100 [=====] - 2s 20ms/step - loss: 4.2123
Epoch 80/100
100/100 [=====] - 3s 27ms/step - loss: 3.8738
Epoch 81/100
100/100 [=====] - 2s 21ms/step - loss: 3.8756
Epoch 82/100
100/100 [=====] - 2s 20ms/step - loss: 3.4062
Epoch 83/100
100/100 [=====] - 2s 20ms/step - loss: 3.7847
Epoch 84/100
100/100 [=====] - 2s 21ms/step - loss: 3.8969
Epoch 85/100
100/100 [=====] - 2s 25ms/step - loss: 3.7772
Epoch 86/100
100/100 [=====] - 2s 23ms/step - loss: 3.7539
Epoch 87/100
100/100 [=====] - 2s 19ms/step - loss: 3.9706
Epoch 88/100
100/100 [=====] - 2s 20ms/step - loss: 3.5051
Epoch 89/100
100/100 [=====] - 2s 20ms/step - loss: 3.4315
Epoch 90/100
100/100 [=====] - 2s 20ms/step - loss: 3.9469
Epoch 91/100
100/100 [=====] - 3s 27ms/step - loss: 3.8497
Epoch 92/100
100/100 [=====] - 2s 21ms/step - loss: 3.8383
Epoch 93/100
100/100 [=====] - 2s 22ms/step - loss: 3.7502
Epoch 94/100
100/100 [=====] - 5s 48ms/step - loss: 3.7458
Epoch 95/100
100/100 [=====] - 4s 38ms/step - loss: 3.9495
Epoch 96/100
100/100 [=====] - 2s 20ms/step - loss: 3.6007
Epoch 97/100
100/100 [=====] - 2s 21ms/step - loss: 3.9899
Epoch 98/100
100/100 [=====] - 5s 45ms/step - loss: 3.5957
Epoch 99/100
100/100 [=====] - 3s 33ms/step - loss: 4.2479
Epoch 100/100
100/100 [=====] - 3s 32ms/step - loss: 3.4564
1/1 [=====] - 1s 521ms/step
```

```
<ipython-input-20-cbba94134022>:107: FutureWarning: `multichannel` is a deprecated argument name for `structural_similarity`. It will be removed in version 1.0. Please use `channel_axis` instead.
```

```
current_ssim_bicubic = ssim(hr_img_np, upscaled_img_bicubic_np, multichannel=True)
```

```
<ipython-input-20-cbba94134022>:112: FutureWarning: `multichannel` is a deprecated argument name for `structural_similarity`. It will be removed in version 1.0. Please use `channel_axis` instead.
```


```
current_ssim_model = ssim(hr_img_np, upscaled_img_model_np, multichannel=True)
```



1/1 [=====] - 0s 20ms/step




1/1 [=====] - 0s 19ms/step




1/1 [=====] - 0s 20ms/step




1/1 [=====] - 0s 33ms/step




1/1 [=====] - 0s 19ms/step




1/1 [=====] - 0s 19ms/step




1/1 [=====] - 0s 19ms/step



1/1 [=====] - 0s 19ms/step



1/1 [=====] - 0s 20ms/step



Average PSNR (Bicubic): 34.97  
 Average SSIM (Bicubic): 0.917  
 Average PSNR (Model): 36.00  
 Average SSIM (Model): 0.937

```
In [22]: import tensorflow as tf
from tensorflow.keras.layers import Input, Conv2D, Conv2DTranspose, LeakyReLU, Lambda, A
from tensorflow.keras.models import Model
from tensorflow.keras.optimizers import Adam
from tensorflow.keras.applications import VGG19
from tensorflow.keras.preprocessing.image import ImageDataGenerator
from tensorflow.keras.optimizers.schedules import ExponentialDecay
import numpy as np
from skimage.metrics import structural_similarity as ssim
from skimage.metrics import peak_signal_noise_ratio as psnr
import matplotlib.pyplot as plt

def create_model1(scale=2):
    inputs = tf.keras.Input(shape=(None, None, 3))
    x = Conv2D(128, (5, 5), padding='same')(inputs)
    x = LeakyReLU(alpha=0.2)(x)

    for _ in range(8):
        x = residual_block(x, 128)

    x = Conv2D(128 * (scale ** 2), (3, 3), padding='same')(x)
    x = tf.nn.depth_to_space(x, scale)
    x = LeakyReLU(alpha=0.2)(x)
    x = Conv2D(3, (5, 5), padding='same')(x)

    outputs = Add()([x, tf.keras.layers.UpSampling2D(size=(scale, scale), interpolation=
    model = Model(inputs, outputs)
    return model

vgg = VGG19(weights='imagenet', include_top=False, input_shape=(None, None, 3))
vgg.trainable = False

def perceptual_loss(y_true, y_pred):
    vgg_true = vgg(y_true)
    vgg_pred = vgg(y_pred)
    return tf.reduce_mean(tf.square(vgg_true - vgg_pred))

model.compile(optimizer=Adam(learning_rate=1e-4), loss=['mean_absolute_error', perce

datagen = ImageDataGenerator(
    rotation_range=20,
    width_shift_range=0.1,
    height_shift_range=0.1,
    horizontal_flip=True
)

lr_schedule = tf.keras.optimizers.schedules.ExponentialDecay(
    initial_learning_rate=1e-4,
```

```

decay_steps=1000,
decay_rate=0.95,
staircase=True
)
optimizer = Adam(learning_rate=lr_schedule)

def channel_attention(x, reduction_ratio=16):
    _, _, _, c = x.shape
    avg_pool = GlobalAveragePooling2D()(x)
    avg_pool = Dense(c // reduction_ratio, activation='relu')(avg_pool)
    avg_pool = Dense(c, activation='sigmoid')(avg_pool)
    return Multiply()([x, avg_pool])

def residual_block(x, filters):
    residual = x
    x = Conv2D(filters, (3, 3), padding='same')(x)
    x = LeakyReLU(alpha=0.2)(x)
    x = Conv2D(filters, (3, 3), padding='same')(x)
    x = channel_attention(x)
    x = Add()([x, residual])
    x = LeakyReLU(alpha=0.2)(x)
    return x

def create_model2(scale=2):
    inputs = tf.keras.Input(shape=(None, None, 3))
    x = Conv2D(64, (5, 5), padding='same')(inputs)
    x = LeakyReLU(alpha=0.2)(x)

    for _ in range(6):
        x = residual_block(x, 64)

    x = Conv2D(64 * (scale ** 2), (3, 3), padding='same')(x)
    x = tf.nn.depth_to_space(x, scale)
    x = LeakyReLU(alpha=0.2)(x)
    x = Conv2D(3, (5, 5), padding='same')(x)

    outputs = Add()([x, tf.keras.layers.UpSampling2D(size=(scale, scale), interpolation='nearest')(x)])
    model = Model(inputs, outputs)
    return model

from tensorflow.keras.models import Sequential

#model 3
def create_model3(scale=2):
    model = Sequential()
    model.add(Conv2D(64, (3, 3), padding='same', input_shape=(None, None, 3)))
    model.add(LeakyReLU(alpha=0.2))
    model.add(Conv2D(64, (3, 3), padding='same'))
    model.add(LeakyReLU(alpha=0.2))
    model.add(Conv2DTranspose(64, (3, 3), strides=(scale, scale), padding='same'))
    model.add(LeakyReLU(alpha=0.2))
    model.add(Conv2D(3, (3, 3), padding='same', activation='tanh'))
    return model

def normalize(image):
    image = tf.cast(image, tf.float32)
    return image / 127.5 - 1

def train_and_evaluate_model(model, model_name, val_dataset, epochs, steps_per_epoch):
    model.fit(val_dataset, epochs=epochs, steps_per_epoch=steps_per_epoch)

    psnr_values = []
    ssim_values = []
    for lr, hr in val_dataset:
        sr = model.predict(lr)
        hr = hr.numpy().squeeze()

```

```

sr = sr.squeeze()

psnr_value = psnr(hr, sr)
ssim_value = ssim(hr, sr, multichannel=True)

psnr_values.append(psnr_value)
ssim_values.append(ssim_value)

print(f"Average PSNR for {model_name}: {np.mean(psnr_values)}")
print(f"Average SSIM for {model_name}: {np.mean(ssim_values)}")

return model

def compare_images(models, val_dataset, num_samples=10):
    psnr_values_bicubic = []
    ssim_values_bicubic = []
    psnr_values_models = {model_name: [] for model_name in models}
    ssim_values_models = {model_name: [] for model_name in models}

    for lr_img, hr_img in val_dataset.take(num_samples):
        bicubic_img = tf.image.resize(lr_img, [hr_img.shape[1], hr_img.shape[2]], method

        if bicubic_img.dtype != tf.uint8:
            bicubic_img = tf.clip_by_value(bicubic_img, 0.0, 255.0)
            bicubic_img = tf.cast(bicubic_img, tf.uint8)

        hr_img_np = hr_img.numpy()[0]
        bicubic_img_np = bicubic_img.numpy()[0]

        current_psnr_bicubic = psnr(hr_img_np, bicubic_img_np, data_range=hr_img_np.max()
        current_ssim_bicubic = ssim(hr_img_np, bicubic_img_np, multichannel=True)
        psnr_values_bicubic.append(current_psnr_bicubic)
        ssim_values_bicubic.append(current_ssim_bicubic)

        plt.figure(figsize=(24, 6))
        plt.subplot(1, len(models) + 3, 1)
        plt.imshow(lr_img.numpy()[0])
        plt.title('Low-Resolution')
        plt.axis('off')

        plt.subplot(1, len(models) + 3, 2)
        plt.imshow(bicubic_img_np)
        plt.title(f'Bicubic\nPSNR: {current_psnr_bicubic:.2f}, SSIM: {current_ssim_bicub
        plt.axis('off')

        for j, (model_name, model) in enumerate(models.items()):
            sr_img = model.predict(lr_img)

            if sr_img.dtype != tf.uint8:
                sr_img = tf.clip_by_value(sr_img, 0.0, 255.0)
                sr_img = tf.cast(sr_img, tf.uint8)

            sr_img_np = sr_img.numpy()[0]

            current_psnr_model = psnr(hr_img_np, sr_img_np, data_range=hr_img_np.max() -
            current_ssim_model = ssim(hr_img_np, sr_img_np, multichannel=True)
            psnr_values_models[model_name].append(current_psnr_model)
            ssim_values_models[model_name].append(current_ssim_model)

            plt.subplot(1, len(models) + 3, j + 3)
            plt.imshow(sr_img_np)
            plt.title(f'{model_name}\nPSNR: {current_psnr_model:.2f}, SSIM: {current_ssi
            plt.axis('off')

        plt.subplot(1, len(models) + 3, len(models) + 3)
        plt.imshow(hr_img_np)

```

```

plt.title('High-Resolution')
plt.axis('off')

plt.tight_layout()
plt.show()

print(f'Average PSNR (Bicubic): {np.mean(psnr_values_bicubic):.2f}')
print(f'Average SSIM (Bicubic): {np.mean(ssim_values_bicubic):.3f}')

for model_name in models:
    print(f'Average PSNR ({model_name}): {np.mean(psnr_values_models[model_name]):.2}')
    print(f'Average SSIM ({model_name}): {np.mean(ssim_values_models[model_name]):.3}')

scale = 2
val_div2k = DIV2K(scale=scale, subset='valid', downgrade='bicubic')
val_dataset = val_div2k.dataset(batch_size=1, repeat_count=1)

# Normalize images for model 3
val_dataset_normalized = val_dataset.map(lambda lr, hr: (normalize(lr), normalize(hr)))

val_dataset = val_dataset.prefetch(tf.data.experimental.AUTOTUNE)
val_dataset_normalized = val_dataset_normalized.prefetch(tf.data.experimental.AUTOTUNE)

total_samples = len(val_div2k)
batch_size = 1
steps_per_epoch = total_samples // batch_size

# Train and evaluate each model
models = {}

# Model 1
model1 = create_model1(scale)
model1.compile(optimizer=Adam(learning_rate=1e-4), loss=['mean_absolute_error', perceptu
models['Model 3'] = train_and_evaluate_model(model1, 'Model 3', val_dataset, epochs=50,

# Model 2
model2 = create_model2(scale)
model2.compile(optimizer=Adam(learning_rate=1e-4), loss='mean_absolute_error')
models['Model 2'] = train_and_evaluate_model(model2, 'Model 2', val_dataset, epochs=50,

# Model 3
model3 = create_model3(scale)
model3.compile(optimizer=Adam(learning_rate=1e-4), loss='mean_squared_error')
models['Model 1'] = train_and_evaluate_model(model3, 'Model 1', val_dataset_normalized,

# Compare images from each model
compare_images(models, val_dataset)

```

```
Epoch 1/50
```

```
100/100 [=====] - 12s 23ms/step - loss: 5.3415
```

```
Epoch 2/50
```

```
100/100 [=====] - 2s 20ms/step - loss: 4.5209
```

```
Epoch 3/50
```

```
100/100 [=====] - 2s 22ms/step - loss: 4.7110
```

```
Epoch 4/50
```

```
100/100 [=====] - 2s 21ms/step - loss: 5.0533
```

```
Epoch 5/50
```

```
100/100 [=====] - 3s 27ms/step - loss: 4.6831
```

```
Epoch 6/50
```

```
100/100 [=====] - 2s 20ms/step - loss: 4.8695
```

```
Epoch 7/50
```

```
100/100 [=====] - 2s 21ms/step - loss: 4.1016
```

```
Epoch 8/50
```

```
100/100 [=====] - 2s 20ms/step - loss: 4.7641
```

```
Epoch 9/50
```

```
100/100 [=====] - 2s 21ms/step - loss: 4.0010
```

Epoch 10/50  
100/100 [=====] - 3s 26ms/step - loss: 4.8281  
Epoch 11/50  
100/100 [=====] - 2s 22ms/step - loss: 4.7039  
Epoch 12/50  
100/100 [=====] - 2s 22ms/step - loss: 4.1234  
Epoch 13/50  
100/100 [=====] - 2s 21ms/step - loss: 4.4925  
Epoch 14/50  
100/100 [=====] - 2s 24ms/step - loss: 3.7893  
Epoch 15/50  
100/100 [=====] - 3s 25ms/step - loss: 4.4041  
Epoch 16/50  
100/100 [=====] - 2s 21ms/step - loss: 4.0777  
Epoch 17/50  
100/100 [=====] - 2s 21ms/step - loss: 5.2249  
Epoch 18/50  
100/100 [=====] - 2s 22ms/step - loss: 4.1420  
Epoch 19/50  
100/100 [=====] - 2s 25ms/step - loss: 4.2884  
Epoch 20/50  
100/100 [=====] - 2s 23ms/step - loss: 4.3409  
Epoch 21/50  
100/100 [=====] - 2s 21ms/step - loss: 4.4569  
Epoch 22/50  
100/100 [=====] - 2s 24ms/step - loss: 4.1945  
Epoch 23/50  
100/100 [=====] - 2s 24ms/step - loss: 4.5564  
Epoch 24/50  
100/100 [=====] - 2s 20ms/step - loss: 4.4726  
Epoch 25/50  
100/100 [=====] - 2s 21ms/step - loss: 3.9612  
Epoch 26/50  
100/100 [=====] - 2s 20ms/step - loss: 3.8823  
Epoch 27/50  
100/100 [=====] - 2s 21ms/step - loss: 3.7541  
Epoch 28/50  
100/100 [=====] - 3s 26ms/step - loss: 5.1146  
Epoch 29/50  
100/100 [=====] - 2s 22ms/step - loss: 4.1593  
Epoch 30/50  
100/100 [=====] - 2s 21ms/step - loss: 4.2379  
Epoch 31/50  
100/100 [=====] - 2s 23ms/step - loss: 3.8970  
Epoch 32/50  
100/100 [=====] - 3s 25ms/step - loss: 4.1728  
Epoch 33/50  
100/100 [=====] - 2s 21ms/step - loss: 3.9019  
Epoch 34/50  
100/100 [=====] - 2s 20ms/step - loss: 3.8172  
Epoch 35/50  
100/100 [=====] - 3s 25ms/step - loss: 4.2406  
Epoch 36/50  
100/100 [=====] - 2s 21ms/step - loss: 4.3094  
Epoch 37/50  
100/100 [=====] - 2s 22ms/step - loss: 3.9017  
Epoch 38/50  
100/100 [=====] - 2s 22ms/step - loss: 3.9764  
Epoch 39/50  
100/100 [=====] - 3s 28ms/step - loss: 3.5434  
Epoch 40/50  
100/100 [=====] - 2s 20ms/step - loss: 4.0876  
Epoch 41/50  
100/100 [=====] - 2s 21ms/step - loss: 3.6392  
Epoch 42/50  
100/100 [=====] - 2s 20ms/step - loss: 3.9227



1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 21ms/step  
1/1 [=====] - 0s 20ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 22ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 32ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 21ms/step  
1/1 [=====] - 0s 21ms/step  
1/1 [=====] - 0s 20ms/step  
1/1 [=====] - 0s 20ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 22ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 20ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 20ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 40ms/step  
1/1 [=====] - 0s 31ms/step  
1/1 [=====] - 0s 33ms/step  
1/1 [=====] - 0s 33ms/step  
1/1 [=====] - 0s 32ms/step  
1/1 [=====] - 0s 34ms/step  
1/1 [=====] - 0s 37ms/step

Average PSNR for Model 3: 36.005349442905505

Average SSIM for Model 3: 0.9280159842924234

Epoch 1/50

100/100 [=====] - 9s 17ms/step - loss: 5.3696

Epoch 2/50

100/100 [=====] - 2s 17ms/step - loss: 4.9430

Epoch 3/50

100/100 [=====] - 2s 20ms/step - loss: 4.8713

Epoch 4/50  
100/100 [=====] - 2s 23ms/step - loss: 4.5624  
Epoch 5/50  
100/100 [=====] - 2s 17ms/step - loss: 4.5936  
Epoch 6/50  
100/100 [=====] - 2s 16ms/step - loss: 4.9759  
Epoch 7/50  
100/100 [=====] - 2s 16ms/step - loss: 4.5547  
Epoch 8/50  
100/100 [=====] - 2s 23ms/step - loss: 5.3057  
Epoch 9/50  
100/100 [=====] - 2s 16ms/step - loss: 4.5518  
Epoch 10/50  
100/100 [=====] - 2s 18ms/step - loss: 4.5135  
Epoch 11/50  
100/100 [=====] - 2s 18ms/step - loss: 4.8147  
Epoch 12/50  
100/100 [=====] - 2s 17ms/step - loss: 4.5370  
Epoch 13/50  
100/100 [=====] - 2s 18ms/step - loss: 4.6813  
Epoch 14/50  
100/100 [=====] - 2s 23ms/step - loss: 4.4705  
Epoch 15/50  
100/100 [=====] - 2s 16ms/step - loss: 4.6332  
Epoch 16/50  
100/100 [=====] - 2s 17ms/step - loss: 3.9183  
Epoch 17/50  
100/100 [=====] - 2s 16ms/step - loss: 4.6467  
Epoch 18/50  
100/100 [=====] - 2s 17ms/step - loss: 4.4595  
Epoch 19/50  
100/100 [=====] - 2s 16ms/step - loss: 4.5318  
Epoch 20/50  
100/100 [=====] - 2s 19ms/step - loss: 4.4685  
Epoch 21/50  
100/100 [=====] - 2s 16ms/step - loss: 4.0639  
Epoch 22/50  
100/100 [=====] - 2s 18ms/step - loss: 4.7729  
Epoch 23/50  
100/100 [=====] - 2s 17ms/step - loss: 4.2310  
Epoch 24/50  
100/100 [=====] - 2s 17ms/step - loss: 4.2481  
Epoch 25/50  
100/100 [=====] - 2s 20ms/step - loss: 4.2921  
Epoch 26/50  
100/100 [=====] - 2s 17ms/step - loss: 4.2914  
Epoch 27/50  
100/100 [=====] - 2s 17ms/step - loss: 4.2811  
Epoch 28/50  
100/100 [=====] - 2s 16ms/step - loss: 4.0340  
Epoch 29/50  
100/100 [=====] - 2s 16ms/step - loss: 4.1832  
Epoch 30/50  
100/100 [=====] - 2s 16ms/step - loss: 4.3801  
Epoch 31/50  
100/100 [=====] - 2s 22ms/step - loss: 4.9571  
Epoch 32/50  
100/100 [=====] - 2s 17ms/step - loss: 4.2957  
Epoch 33/50  
100/100 [=====] - 2s 17ms/step - loss: 3.7843  
Epoch 34/50  
100/100 [=====] - 2s 18ms/step - loss: 4.2908  
Epoch 35/50  
100/100 [=====] - 2s 17ms/step - loss: 4.1710  
Epoch 36/50  
100/100 [=====] - 2s 17ms/step - loss: 4.1996

```
Epoch 37/50
100/100 [=====] - 2s 23ms/step - loss: 4.2383
Epoch 38/50
100/100 [=====] - 2s 16ms/step - loss: 3.9751
Epoch 39/50
100/100 [=====] - 2s 16ms/step - loss: 4.6707
Epoch 40/50
100/100 [=====] - 2s 17ms/step - loss: 4.1028
Epoch 41/50
100/100 [=====] - 2s 17ms/step - loss: 4.2331
Epoch 42/50
100/100 [=====] - 2s 23ms/step - loss: 4.3261
Epoch 43/50
100/100 [=====] - 2s 18ms/step - loss: 4.0355
Epoch 44/50
100/100 [=====] - 2s 17ms/step - loss: 4.2321
Epoch 45/50
100/100 [=====] - 2s 18ms/step - loss: 4.1401
Epoch 46/50
100/100 [=====] - 2s 18ms/step - loss: 3.7876
Epoch 47/50
100/100 [=====] - 2s 17ms/step - loss: 4.1154
Epoch 48/50
100/100 [=====] - 2s 20ms/step - loss: 4.2950
Epoch 49/50
100/100 [=====] - 2s 18ms/step - loss: 3.8308
Epoch 50/50
100/100 [=====] - 2s 17ms/step - loss: 4.7636
1/1 [=====] - 0s 271ms/step
1/1 [=====] - 0s 18ms/step
1/1 [=====] - 0s 18ms/step
1/1 [=====] - 0s 18ms/step
1/1 [=====] - 0s 19ms/step
1/1 [=====] - 0s 19ms/step
1/1 [=====] - 0s 18ms/step
1/1 [=====] - 0s 19ms/step
1/1 [=====] - 0s 18ms/step
1/1 [=====] - 0s 19ms/step
1/1 [=====] - 0s 18ms/step
1/1 [=====] - 0s 17ms/step
1/1 [=====] - 0s 18ms/step
1/1 [=====] - 0s 19ms/step
1/1 [=====] - 0s 18ms/step
1/1 [=====] - 0s 18ms/step
1/1 [=====] - 0s 30ms/step
1/1 [=====] - 0s 20ms/step
1/1 [=====] - 0s 18ms/step
1/1 [=====] - 0s 19ms/step
1/1 [=====] - 0s 20ms/step
1/1 [=====] - 0s 18ms/step
1/1 [=====] - 0s 20ms/step
1/1 [=====] - 0s 22ms/step
1/1 [=====] - 0s 19ms/step
1/1 [=====] - 0s 21ms/step
1/1 [=====] - 0s 19ms/step
1/1 [=====] - 0s 21ms/step
1/1 [=====] - 0s 22ms/step
1/1 [=====] - 0s 20ms/step
1/1 [=====] - 0s 18ms/step
1/1 [=====] - 0s 19ms/step
1/1 [=====] - 0s 19ms/step
1/1 [=====] - 0s 20ms/step
1/1 [=====] - 0s 19ms/step
1/1 [=====] - 0s 22ms/step
1/1 [=====] - 0s 18ms/step
1/1 [=====] - 0s 18ms/step
```

1/1 [=====] - 0s 21ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 20ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 20ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 22ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 20ms/step  
1/1 [=====] - 0s 21ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 26ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 34ms/step  
1/1 [=====] - 0s 32ms/step  
1/1 [=====] - 0s 29ms/step  
1/1 [=====] - 0s 28ms/step  
1/1 [=====] - 0s 31ms/step  
1/1 [=====] - 0s 28ms/step  
1/1 [=====] - 0s 27ms/step  
1/1 [=====] - 0s 27ms/step  
1/1 [=====] - 0s 28ms/step  
1/1 [=====] - 0s 27ms/step  
1/1 [=====] - 0s 32ms/step  
1/1 [=====] - 0s 30ms/step  
1/1 [=====] - 0s 28ms/step  
1/1 [=====] - 0s 33ms/step  
1/1 [=====] - 0s 29ms/step  
1/1 [=====] - 0s 33ms/step  
1/1 [=====] - 0s 33ms/step

Average PSNR for Model 2: 36.13714561155611

Average SSIM for Model 2: 0.9212456011135979

Epoch 1/50

100/100 [=====] - 3s 6ms/step - loss: 0.1458

Epoch 2/50  
100/100 [=====] - 1s 7ms/step - loss: 0.0485  
Epoch 3/50  
100/100 [=====] - 1s 7ms/step - loss: 0.0325  
Epoch 4/50  
100/100 [=====] - 1s 6ms/step - loss: 0.0276  
Epoch 5/50  
100/100 [=====] - 1s 6ms/step - loss: 0.0230  
Epoch 6/50  
100/100 [=====] - 1s 6ms/step - loss: 0.0204  
Epoch 7/50  
100/100 [=====] - 1s 6ms/step - loss: 0.0163  
Epoch 8/50  
100/100 [=====] - 1s 10ms/step - loss: 0.0181  
Epoch 9/50  
100/100 [=====] - 1s 7ms/step - loss: 0.0160  
Epoch 10/50  
100/100 [=====] - 1s 6ms/step - loss: 0.0140  
Epoch 11/50  
100/100 [=====] - 1s 6ms/step - loss: 0.0135  
Epoch 12/50  
100/100 [=====] - 1s 6ms/step - loss: 0.0126  
Epoch 13/50  
100/100 [=====] - 1s 6ms/step - loss: 0.0121  
Epoch 14/50  
100/100 [=====] - 1s 6ms/step - loss: 0.0135  
Epoch 15/50  
100/100 [=====] - 1s 6ms/step - loss: 0.0103  
Epoch 16/50  
100/100 [=====] - 1s 6ms/step - loss: 0.0127  
Epoch 17/50  
100/100 [=====] - 1s 6ms/step - loss: 0.0114  
Epoch 18/50  
100/100 [=====] - 1s 6ms/step - loss: 0.0098  
Epoch 19/50  
100/100 [=====] - 1s 6ms/step - loss: 0.0099  
Epoch 20/50  
100/100 [=====] - 1s 6ms/step - loss: 0.0108  
Epoch 21/50  
100/100 [=====] - 1s 10ms/step - loss: 0.0101  
Epoch 22/50  
100/100 [=====] - 1s 7ms/step - loss: 0.0099  
Epoch 23/50  
100/100 [=====] - 1s 6ms/step - loss: 0.0090  
Epoch 24/50  
100/100 [=====] - 1s 6ms/step - loss: 0.0093  
Epoch 25/50  
100/100 [=====] - 1s 6ms/step - loss: 0.0086  
Epoch 26/50  
100/100 [=====] - 1s 7ms/step - loss: 0.0117  
Epoch 27/50  
100/100 [=====] - 1s 7ms/step - loss: 0.0098  
Epoch 28/50  
100/100 [=====] - 1s 7ms/step - loss: 0.0091  
Epoch 29/50  
100/100 [=====] - 1s 6ms/step - loss: 0.0087  
Epoch 30/50  
100/100 [=====] - 1s 6ms/step - loss: 0.0092  
Epoch 31/50  
100/100 [=====] - 1s 7ms/step - loss: 0.0080  
Epoch 32/50  
100/100 [=====] - 1s 7ms/step - loss: 0.0084  
Epoch 33/50  
100/100 [=====] - 1s 11ms/step - loss: 0.0078  
Epoch 34/50  
100/100 [=====] - 1s 6ms/step - loss: 0.0092



1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 17ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 21ms/step  
1/1 [=====] - 0s 16ms/step  
1/1 [=====] - 0s 16ms/step  
1/1 [=====] - 0s 16ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 17ms/step  
1/1 [=====] - 0s 17ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 17ms/step  
1/1 [=====] - 0s 16ms/step  
1/1 [=====] - 0s 16ms/step  
1/1 [=====] - 0s 16ms/step  
1/1 [=====] - 0s 24ms/step  
1/1 [=====] - 0s 17ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 23ms/step  
1/1 [=====] - 0s 26ms/step  
1/1 [=====] - 0s 29ms/step  
1/1 [=====] - 0s 23ms/step  
1/1 [=====] - 0s 29ms/step  
1/1 [=====] - 0s 24ms/step  
1/1 [=====] - 0s 26ms/step  
1/1 [=====] - 0s 28ms/step  
1/1 [=====] - 0s 29ms/step  
1/1 [=====] - 0s 24ms/step  
1/1 [=====] - 0s 24ms/step  
1/1 [=====] - 0s 27ms/step  
1/1 [=====] - 0s 26ms/step  
1/1 [=====] - 0s 24ms/step  
1/1 [=====] - 0s 32ms/step  
1/1 [=====] - 0s 24ms/step  
1/1 [=====] - 0s 25ms/step  
1/1 [=====] - 0s 24ms/step  
1/1 [=====] - 0s 27ms/step  
1/1 [=====] - 0s 28ms/step  
1/1 [=====] - 0s 34ms/step  
1/1 [=====] - 0s 26ms/step  
1/1 [=====] - 0s 29ms/step  
1/1 [=====] - 0s 16ms/step  
1/1 [=====] - 0s 17ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 17ms/step  
1/1 [=====] - 0s 16ms/step  
1/1 [=====] - 0s 19ms/step  
1/1 [=====] - 0s 17ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 24ms/step  
1/1 [=====] - 0s 18ms/step  
1/1 [=====] - 0s 17ms/step  
1/1 [=====] - 0s 23ms/step  
1/1 [=====] - 0s 16ms/step  
1/1 [=====] - 0s 16ms/step

Average PSNR for Model 1: 28.315669545395334

Average SSIM for Model 1: 0.863423764705658

1/1 [=====] - 0s 21ms/step

1/1 [=====] - ETA: 0s

<ipython-input-22-1682c2cff82a>:145: FutureWarning: `multichannel` is a deprecated argument name for `structural\_similarity`. It will be removed in version 1.0. Please use `channel\_axis` instead.

current\_ssim\_bicubic = ssim(hr\_img\_np, bicubic\_img\_np, multichannel=True)

<ipython-input-22-1682c2cff82a>:170: FutureWarning: `multichannel` is a deprecated argument name for `structural\_similarity`. It will be removed in version 1.0. Please use `channel\_axis` instead.

current\_ssim\_model = ssim(hr\_img\_np, sr\_img\_np, multichannel=True)

1/1 [=====] - 0s 19ms/step

1/1 [=====] - 0s 75ms/step



1/1 [=====] - 0s 20ms/step

1/1 [=====] - 0s 19ms/step

1/1 [=====] - 0s 17ms/step



1/1 [=====] - 0s 22ms/step

1/1 [=====] - 0s 21ms/step

1/1 [=====] - 0s 19ms/step



1/1 [=====] - 0s 22ms/step

1/1 [=====] - 0s 21ms/step


1/1 [=====] - 0s 24ms/step




1/1 [=====] - 0s 27ms/step

1/1 [=====] - 0s 21ms/step


1/1 [=====] - 0s 19ms/step




1/1 [=====] - 0s 21ms/step  
 1/1 [=====] - 0s 20ms/step  
 1/1 [=====] - 0s 18ms/step




1/1 [=====] - 0s 21ms/step  
 1/1 [=====] - 0s 27ms/step  
 1/1 [=====] - 0s 18ms/step




1/1 [=====] - 0s 21ms/step  
 1/1 [=====] - 0s 19ms/step  
 1/1 [=====] - 0s 19ms/step



1/1 [=====] - 0s 20ms/step  
 1/1 [=====] - 0s 19ms/step  
 1/1 [=====] - 0s 17ms/step



1/1 [=====] - 0s 21ms/step  
 1/1 [=====] - 0s 20ms/step  
 1/1 [=====] - 0s 18ms/step



Average PSNR (Bicubic): 32.29

```
Average SSIM (Bicubic): 0.904
Average PSNR (Model 3): 34.00
Average SSIM (Model 3): 0.927
Average PSNR (Model 2): 33.71
Average SSIM (Model 2): 0.924
Average PSNR (Model 1): 4.80
Average SSIM (Model 1): 0.025
```

In [ ]: