ORIGINAL RESEARCH

# Restoring Resolution in Accelerated FLAIR MRI via U-Net Super-Resolution

**Eugene Joo '27**
Washington University in St. Louis

**Accelerated MRI scanning often results in lower resolution, introducing image distortions and loss of critical details. This study explores the application of machine learning to super-resolve Fluid-Attenuated Inversion Recovery Magnetic Resonance Imaging (FLAIR MRI) scans and restore those missing details. Specifically, we test models based on the U-Net architecture trained with two down sampling methods: removing two out of every 3 slices and linearly interpolating in image space and cropping the top and bottom thirds slices in k-space, the raw frequency data. Our experiments demonstrate that k-space down sampling consistently outperforms image-space methods in both Structural Similarity Index (SSIM) and Peak Signal-to-Noise Ratio (PSNR) metrics, while also achieving superior reconstruction of fine details. Additionally, a multi-slice input approach using three adjacent slices was shown to further improve results by providing additional spatial context. This work allows for the restoration of MRI scans collected in an accelerated manner, significantly enhancing image quality and detail for improved diagnostic confidence.**

Magnetic resonance imaging (MRI) is a critical method in the non-invasive detection and diagnosis of degenerative diseases in the brain. The precision of a diagnosis using MRI scans is highly dependent on their spatial resolution, with higher spatial resolution providing a greater level of detail in the final image. However, achieving this higher resolution requires prolonged scan times, which is sometimes not an option in time-sensitive scenarios.

Scan times can be reduced by collecting less data in k-space, which is the frequency domain where raw MRI data is collected and stored [1]. However, this reduction in data lowers the spatial resolution and can lead to a low-resolution final image containing artifacts such as blurring.

This work attempts to address these issues by applying machine learning to enhance and super-resolve the MRI images after their acquisition. Specifically, we will focus on super-resolving Fluid Attenuated Inversion Recovery (FLAIR) MRI scans using the U-Net architecture.

Originally developed for biomedical image segmentation, the process of identifying and labeling structures within medical images, U-Net's encoder–decoder structure with skip connections makes it well suited for reconstructing high-resolution images from low-resolution inputs (Figure 1). This capability makes it an ideal choice for restoring fine structural details lost during accelerated scanning [3].

The FLAIR MRI dataset used in this study consists of volumetric scans, which are each stored as a three-dimensional array of voxels (three-dimension pixels). Because the data was collected using an accelerated protocol, two out of every three slices are missing, meaning the images retain only one-third of the original spatial information.

## Methods

To reconstruct the missing information in super-resolution, a model needs to be trained using a simulated data set created from MRI scans at the desired resolution. We experimented with two downsampling methods: one in image space and one in k-space.

The image space downsampling is performed by removing two out of every three slices using linear interpolation on the original high-resolution MRI scan to match the dimensions of the real data. Then, to feed the image into the model, every voxel was tripled, matching the dimensions of the ground truth image. Since the real data was already missing two out of every three slices, feeding it into the model only required each of the voxels to be tripled along the axis with missing data.

The k-space downsampling for the simulated data set involved taking the high-resolution image and converting the image to k-space using a Fourier Transform [1]. By removing information from the domain where raw MRI data is initially collected, we hope to create a more accurate simulated downsampling process. Then, to match the amount of missing information from the real data, the first and last third of the k-space images were masked out. Afterward, an inverse Fourier Transform was applied for the final downsampled image. To replicate this process on the real data, the k-space image in the frequency domain was padded on the left and right sides with rectangles of equal dimensions to the original image to triple the width. After an inverse Fourier Transform, an image that matches the dimensions of the simulated data set was obtained. This dimension matching step is critical in getting the expected behavior from the super-resolution model that was trained on the simulated data set.

Since MRI scans include 3D data, two separate models were trained for two different views of the volume: coronal and sagittal, the terms used to describe different planes of the brain. This approach allowed the model to learn patterns specific to each view, improving reconstruction accuracy.

We were also able to take advantage of the data's 3D nature by implementing a modified version of the U-Net model. This worked by allowing the initial downsampling block to accept multiple input channels, each of which was fed with adjacent slices of the MRI. Then, the output of the final upblock was validated against the ground truth of the middle of the input slices. To ensure

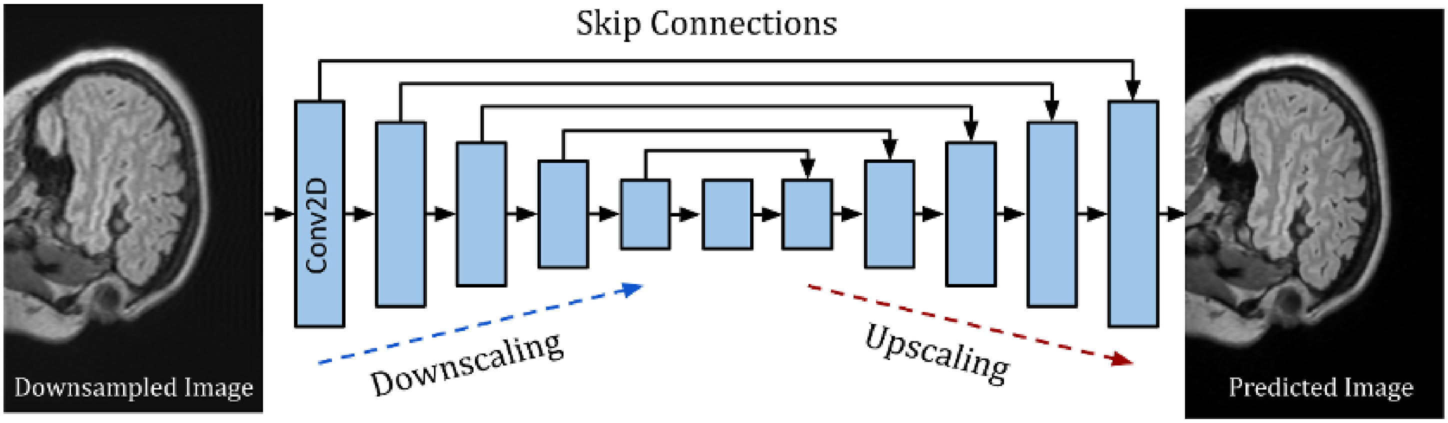


**Fig. 1 |** Diagram of the U-Net architecture used in the model.

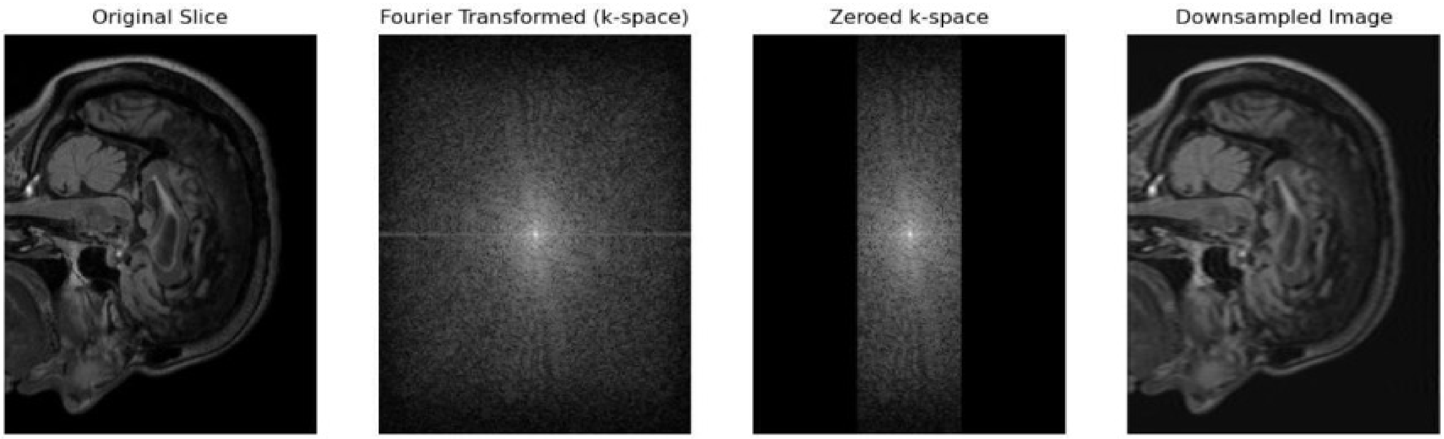


**Fig. 2 |** An example of the k-space downsampling performed for the simulated dataset.

we found the correct balance between giving the model critical extra information and distracting it with irrelevant data, versions of the model accepting one, three, and five input slices were attempted.

Because our model was trained on multiple MRI slices, voxel intensities were normalized to a range of 0–1. This normalization helps the model focus on structural details rather than variations in intensity.

To evaluate the model's performance, we calculated two common image quality metrics: structural similarity index (SSIM) and peak signal-to-noise ratio (PSNR). SSIM measures the perceptual similarity between two images, while PSNR measures how much signal there is compared to noise. For both SSIM and PSNR, a higher value indicates better image quality.

Both metrics were evaluated for input (downsampled and ground truth) and output (predicted and ground truth) and averaged across the validation dataset.

## Results

When applied to their respective downsampling methods, the k-space model consistently outperformed the image-space model in terms of SSIM and PSNR (Table 1). Visual inspection shows that the K-space model preserves finer anatomical details than the image space model (Fig. 3). Thin structures are notably degraded or lost in the image-space model, likely due to the greater information loss introduced by its downsampling method.

We also evaluated configurations of 1, 3, or 5 adjacent MRI slices in addition to the downsampling comparison. Among the tested slice configurations, the 3-slice model achieved the highest SSIM and PSNR scores, outperforming both the 1-slice and 5-slice models. These findings suggest that the 3-slice configuration strikes a balance between providing sufficient contextual information and avoiding overfitting. The 3-slice configuration likely benefits from incorporating sufficient spatial context without overwhelming the model, which might occur with the 5-slice variant. Conversely, the 1-slice model likely lacks adequate contextual information for robust reconstruction.

The models show promising results with both simulated and real-world data. Specifically, the k-space model shows qualitatively better performance when applied to real MRI data. This performance disparity suggests that the k-space approach generalizes more effectively to natural noise and variations found in real-world scenarios, whereas the image space super-resolution process appears more sensitive to these inconsistencies. However, the distribution of the super-resolved data still closely resembles that of the low-resolution data instead of the high-resolution data we are attempting to match (Fig. 6). Additionally, the distribution of the k-space downsampled data more closely matches the high-resolution distribution, revealing an imperfect simulation of the real data.

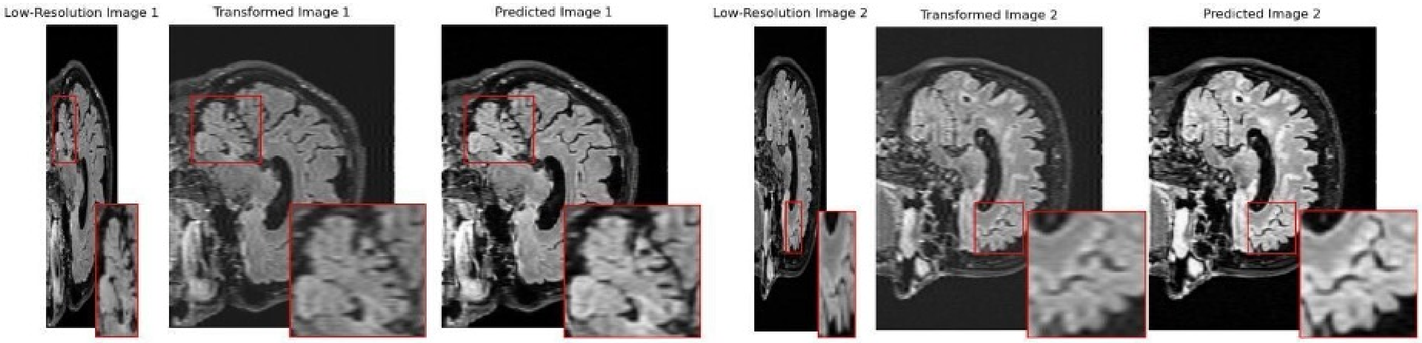


**Fig. 3 |** Two examples of the k-space model with three input slices applied to the real data set. As can be seen in the highlighted areas, the model is qualitatively able to sharpen the images substantially.

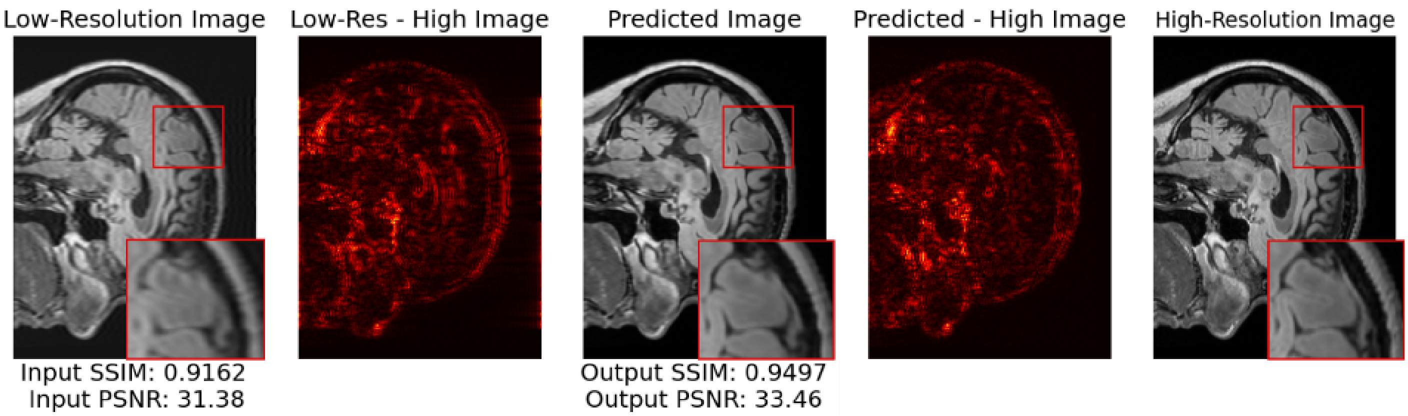


**Fig. 4 |** An example of the output of a model trained with three input channels and data downsampled with the k-space method.

| | Image Space Downsampling **SSIM** | Image Space Downsampling **PSNR** | K-space Downsampling **SSIM** | K-space Downsampling **PSNR** |
|---|---|---|---|---|
| **Input** | 0.8367 | 29.0960 | 0.9420 | 35.1550 |
| **Image Space Model** | 0.9168 | 33.5757 | 0.8856 | 31.3039 |
| **K-space One Slice Model** | 0.8056 | 27.0096 | 0.9580 | 37.0127 |
| **K-space Three Slice Model** | 0.8277 | 28.1409 | **0.9640** | **37.2131** |
| **K-space Five Slice Model** | 0.8213 | 28.0856 | 0.9633 | 37.1335 |

**Table 1 |** Comparison of SSIM and PSNR for all variations of models trained as applied to both image space downsampling done with linear interpolation and k-space downsampling. Input metrics quantify the difference between the downsampled data and the original high-resolution data. For all metrics, a higher value is better.

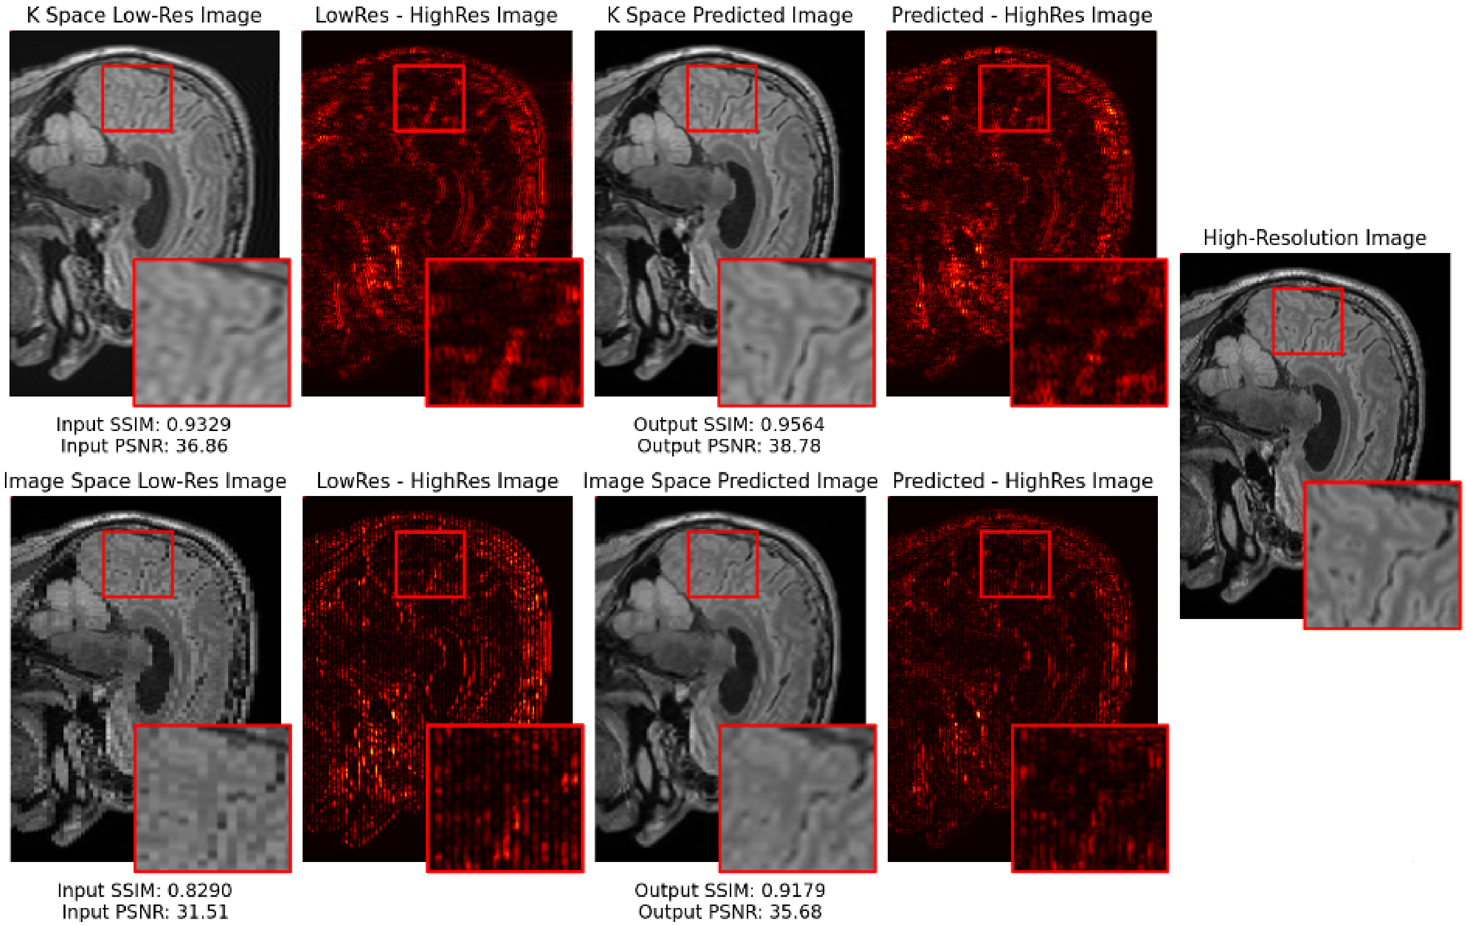


**Fig. 5 |** Comparison between the performance of the K-space and Image-Space models. The K-space model (top) performs much better and can create smooth lines like the ground truth. Meanwhile, the Image-Space model, as featured on the bottom, still carries some artifacts from the downsampling method.

## Conclusions

This project demonstrates the use of the U-Net model architecture, a convolutional neural network (CNN) designed for fast and precise image segmentation, to restore image data lost from accelerated MRI scanning. By comparing downsampling methods performed in image space and k-space, we found that k-space downsampling consistently outperforms image space downsampling in both quantitative metrics (SSIM and PSNR) and qualitative reconstruction of fine anatomical details. Furthermore, experiments with different slice configurations revealed that a 3-slice input strikes an optimal balance between spatial context and model complexity, outperforming 1-slice and 5-slice configurations.

While we were able to achieve relatively high quantitative metrics, the model could be improved by better utilizing the k-space data through deep unfolding, which could be used for additional validation during training [4]. Another limitation lies in the reliance on simulated downsampling, which may not fully reflect real-world MRI conditions. Analysis of the super-resolved data reveals that the distribution still matches the low-resolution data, rather than the desired high-resolution data. While qualitative analysis is still possible with our super-resolved data, this distribution issue could cause problems for certain quantitative pipelines where data fidelity is crucial. Future experiments should be conducted to refine the simulated downsampling process in order to improve our results.

## Acknowledgments

We thank Edward Chandler, Dr. Ulugbek Kamilov, and the rest of the Washington University Computational Imaging Group for providing mentorship and resources essential to this project. We also thank Dr. Matthew R. Brier and Bradley Judge at the Washington University School of Medicine for providing the data that the project was based on.

## References

[1] Allen D. Elster. What is k-space?, 2024.

[2] Michael Lustig, David L. Donoho, and John M. Pauly. Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007.

[3] Philipp Fischer Olaf Ronneberger and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. *Medical Image Computing and Computer-Assisted Intervention* (MICCAI), 9351:234–241, 2015.

[4] Kai Zhang, Luc Van Gool, and Radu Timofte. Deep unfolding network for image super-resolution, 2020.