

**UCLA**

**UCLA Previously Published Works**

**Title**

MRI Super-Resolution with GAN and 3D Multi-Level DenseNet: Smaller, Faster, and Better

**Permalink**

<https://escholarship.org/uc/item/3jv4928r>

**Authors**

Chen, Yuhua  
Christodoulou, Anthony G  
Zhou, Zhengwei  
et al.

**Publication Date**

2020-03-02

# MRI Super-Resolution with GAN and 3D Multi-Level DenseNet: Smaller, Faster, and Better

Yuhua Chen<sup>a,b</sup>, Anthony G. Christodoulou<sup>b</sup>, Zhengwei Zhou<sup>b</sup>, Feng Shi<sup>b</sup>,  
Yibin Xie<sup>b</sup>, Debiao Li<sup>a,b,\*</sup>

<sup>a</sup>*Department of Bioengineering, University of California, Los Angeles, CA, U.S.A*

<sup>b</sup>*Biomedical Imaging Research Institute, Cedars-Sinai Medical Center, CA, U.S.A*

---

## Abstract

High-resolution (HR) magnetic resonance imaging (MRI) provides detailed anatomical information that is critical for diagnosis in the clinical application. However, HR MRI typically comes at the cost of long scan time, small spatial coverage, and low signal-to-noise ratio (SNR). Recent studies showed that with a deep convolutional neural network (CNN), HR generic images could be recovered from low-resolution (LR) inputs via single image super-resolution (SISR) approaches. Additionally, previous works have shown that a deep 3D CNN can generate high-quality SR MRIs by using learned image priors. However, 3D CNN with deep structures, have a large number of parameters and are computationally expensive. In this paper, we propose a novel 3D CNN architecture, namely a multi-level densely connected super-resolution network (mDCSRN), which is light-weight, fast and accurate. We also show that with generative adversarial network (GAN)-guided training, the mDCSRN-GAN provides appealing sharp SR images with rich texture details that are highly comparable with the referenced HR images. Our results from experiments on a large public dataset with 1,113 subjects showed that this new architecture outperformed other popular deep learning methods in recovering 4x resolution-downgraded images in both quality and speed.

*Keywords:* Super-Resolution, Deep Learning, 3D Convolutional Neural Network, MRI

---

## 1. Introduction

High spatial resolution MRI provides important structural details for clinicians to detect disease and make better diagnostic decisions (Pruessner et al., 2000). It provides accurate tissue and organ measurements that benefit quantitative image analysis for better diagnosis and therapeutic monitoring (Greenspan,

---

\*Corresponding author

*Email addresses:* [chyuhua@ucla.edu](mailto:chyuhua@ucla.edu) (Yuhua Chen), [debiao.li@cshs.org](mailto:debiao.li@cshs.org) (Debiao Li)

2008; Park et al., 2003; Xie et al., 2016). However, limited by hardware capacity and patient cooperation, HR imaging is burdened by long scan time, small spatial coverage, and low signal-to-noise ratio (SNR) (Shi et al., 2015). HR MRI is also susceptible to respiratory or internal organ motion (Pang et al., 2016; Zhou et al., 2017), thus it is very difficult if not impossible to perform on moving part of the body (Stucht et al., 2015; Yang et al., 2016). In MRI, the duration between phase encodes is the most time-consuming part of the acquisition process, so scan time increases as spatial resolution improves along phase-encoded dimensions. For example, A 4x resolution-degraded LR MRI would be 4x faster than full-resolution HR, at the cost of losing fine local details. Therefore, with the capability to restore resolution loss in HR from just a single LR image, Single Image Super-Resolution (SISR) (Glasner et al., 2009) is an appealing approach as it promises a reconstructed HR image without adding extra scans or additional multi-image combination processing.

However, the SISR problem is very challenging. Since multiple HR images can be resolution-degraded to the same LR image, SISR is an ill-posed inverse problem. To correctly recover high-frequency details such as local textures and edges from its LR counterparts, an intricate image prior is essential. Previous SISR approaches focus on creating a convex optimization process to find the most likely mapping between LR and HR images (Shi et al., 2015). Constraints are usually applied to regularize such processes. However, the prior knowledge presumed by those constraints does not always hold. One of the popular regularization methods, total variation (Rudin et al., 1992), assumes that the HR image is constant in a small neighborhood, which usually violates the fact that the HR image often carries rich local details and tiny structures, such as intracranial vessels in brain MRI.

In 2D generic images, Dong et al. (2016a,b) show that by utilizing a CNN, the SISR puzzles can be solved with an end-to-end learning-based method. Though a larger neural network with more capacity could help improve the overall performance (Sun et al., 2016), training such a deep CNN has been proven to be difficult (Glorot and Bengio, 2010). Recently, with skip connections (He et al., 2016; Srivastava et al., 2015), embedding (Szegedy et al., 2017), and normalization (Ioffe and Szegedy, 2015), effective training for a deep neural network is now made possible. Kim et al. (2016) showed that a deeper network using all these advanced techniques could achieve significant improvement in SR image quality, showing that the CNN’s architecture is the key to obtain high-quality SR outputs. However, as the network grows deeper, the high-level portion of the network is less likely to make full use of the low-level features due to the vanishing gradient phenomenon (He et al., 2016). Residual learning via skip connection (Ledig et al., 2017) helps to ease the effect. Later, Huang et al. (2017) proved that directly stacking all inputs with CNN feature maps strengthens the information flow, and further reduces gradient vanishing. Additionally, these concatenated layers share features more efficiently, lessen the requirement for the immense amount of parameters usually found in deep neural networks. Hence, Densely connected network (DenseNet), can outperform deep CNNs despite its lighter weight. In SISR, Tong et al. (2017) proposed

SRDenseNet which combines different hierarchy level features into the final reconstruction layer. Their work demonstrates a significant improvement over networks only using high-level features, indicating that multi-level feature fusion is indeed beneficial for the SISR problem. However, there is still room for SRDenseNet to improve, as we will show in the later section.

Following the wave of the rapid progress in natural images, SISR has also been adapted into medical image fields (Litjens et al., 2017; Oktay et al., 2016; You et al., 2019). Most of the existing studies directly borrow the 2D network structure and apply it to medical images slice by slice (Oktay et al., 2016; Wang et al., 2016). However, medical images like Computed Tomography (CT), MRI, and Positron Emission Tomography (PET), often carry anatomy information in 3D. To fully resolve the ill-posed SR problem, a 3D model is more natural and preferable as it can directly extract 3D structural information. Recent studies (Chen et al., 2018; Pham et al., 2017) show that in brain MRI SR, a 3D CNN outperforms its 2D counterpart by a large margin. However, due to the extra dimension introduced by 3D CNN, the parameter number of a deep model also grows at a staggering rate, the so-called curse of dimensionality. For example, a 3D Fast Super-Resolution CNN (FSRCNN) (Dong et al., 2016b) has 5x parameters than a 2D FSRCNN. Almost all recent SISR methods obtain improved performance by adding more weights and layers (Lim et al., 2017; Tai et al., 2017). However, borrowing such idea to 3D is not ideal. An over-parameterized 3D model is much more heavily weighted, computationally expensive, and less practical with the potential of exceeding the computer’s memory limitation. Besides, most of the previous CNN SISR approaches are optimized by the pixel/voxel-wise rectilinear or Euclidean distance ( $L_1/L_2$  loss) between model output and ground truth image. As noticed in Ledig et al. (2017), this loss and its derived Peak Signal to Noise Ratio (PSNR) cannot accurately reflect the perceptual quality of the reconstructed image (Johnson et al., 2016). Therefore, merely taking account of the intensity difference results in suboptimal fuzzy output.

In this paper, we propose a 3D Multi-Level Densely Connected Super-Resolution Network (mDCSRN) and mDCSRN-GAN with an adversarial loss guided training. Our goal is to build a small, fast, but accurate network structure for the SISR system that can recover 3D details from resolution-reduced MRI. We first experimented with our mDCSRN with  $L_1$  loss. Measured by numeric metrics, our mDCSRN outperformed interpolation and popular neural networks while using minimal computational resources. Then we experimented that when trained with a Generative Adversarial Network (GAN) (Goodfellow et al., 2014), our mDCSRN-GAN provided even sharper and abundant detailed texture SR images that are highly comparable with the HR images.

We summarize four main contributions of this work:

- We proposed a 3D multi-level densely connected super-resolution neural network (mDCSRN) which has multi-level direct access to all former image features. It is efficient in memory usage yet provides high-quality SR images, making it practical for 3D medical image data.


- We proposed a bottleneck compressor module with a fixed-number width before each DenseBlock, which helps balance the layer size in different conceptual levels. The compressor greatly reduces memory usage and increases runtime speed without sacrificing performance.
- We proposed a direct combination mechanism that actively feeds all levels’ image feature to the final output. This design enables unobstructed gradient flow for easier training and faster convergence. It also makes use of the effect of model ensemble, further boosting performance.
- We proposed an mDCSRN-GAN that can produce accurate and realistic-looking SR images by applying a 3D generative adversarial network (GAN) during training. Testing on real-world data showed that our GAN network is robust across different platforms and scanners.

## 2. Related Work

**Single Image Super-Resolution.** As a classic problem in computer vision, SISR has been studied for decades. Before deep learning approaches dominated the state-of-the-art performance, SISR techniques mostly relied on interpolation, edge-preservation, statistical analysis, and sparse dictionary learning, which have been well-summarized by Yang et al. (2014). Dong et al. (2016a) were the first to propose a SISR based on a three-layer CNN. They showed that a neural network, namely a Super-Resolution Convolutional Neural Network (SRCNN), is naturally capable of handling feature extraction, feature space building, and image reconstruction together through end-to-end training. SRCNN and its recent version Fast SRCNN (FSRCNN) achieved remarkable performance. Their work has inspired many follow-up studies with more advanced network structures (Kim et al., 2016; Lim et al., 2017; Tai et al., 2017; Tong et al., 2017).

**Efficient Network with Skip Connections.** The performance of the deep learning model keeps improving. However, most of the achievement is built upon the significantly increased model size, wherein the depth of the network becomes a practical issue. As the back-propagated gradients often vanish in the long pathway, it is unlikely to train very deep CNNs. To address this problem, Srivastava et al. (2015) (Highway Network) and He et al. (2016) (ResiNet) proposed the bypassing path, or the skip connection, to add the previous layer to the next for smoother information flows. Huang et al. (2017) discovered that by concatenating previous layers, the network is more efficient and outperforms ResiNet with less number of parameters. As all segments in a DenseNet are directly linked, the gradient can flow unobstructed. Additionally, the dense connections encourage layers to share their features. It dramatically reduces the number of parameters, making the model computational efficient, more robust to new data, and faster to converge.

**Super Resolution with Perceptual Loss.** The most straightforward objective for a super-resolution model to optimize would be the voxel-wise difference between model output and the ground-truth image like  $L_1$  or  $L_2$  loss.



**Fig. 1.** Visual quality comparison between Nearest Neighbor Interpolation, deep neural network optimized for intensity difference, deep neural network optimized for a loss with perceptual penalty, and original HR image with PSNR and SSIM shown above the images. (2 x 2 x 1 resolution degrading)

However, this difference only takes account of the intensity values' dissimilarity between the reconstructed image and the original image, but not the visual quality which more focuses on sharpness and validity of restored structures. Optimizing the voxel-wise difference will force the model to stack and average all the possible HR candidates in image space. Since a voxel-wise loss doesn't account for the perceptual level of information, despite its results have less intensity error on average, it provides over-blurred and implausible results for the human eye. Therefore, as shown in **Fig. 1**, though the voxel-wise loss guided SR model provides a better score in PSNR and SSIM, the model with perceptual loss estimated by a Generative Adversarial Network (GAN) provides more realistic-looking images.

### 3. Method

We designed our SISR model to learn an accurate inverse mapping of the LR image to the reference HR image during an end-to-end training process. The network is fed with LR images, and it outputs resolution-restored SR images. The HR images were only used in training as the target for the system to optimize. A loss function calculated from SR and HR is back-propagated through layers to adjust weights during training. In the deployment phase, the model only reads LR images and produces SR outputs. We will detail our proposed mDCSRN and the GAN-guided training process in the following sections.

### 3.1. SISR Background

A SISR system is a feed-forward model to transform an LR image  $Y$  into an HR image  $X$ . A mathematical representation of the resolution downgrading process from  $X$  to  $Y$  can be written as:

$$Y = f(X), \quad (1)$$

where  $f$  is an arbitrary continuous or discrete function that causes the resolution loss. The SR process is to find an optimal inverse mapping function  $g(\cdot) \approx f^{-1}(\cdot)$ , where  $f^{-1}$  represents the inverse of  $f$ . The recovered HR image, or SR,  $\tilde{X}$  will be:


$$\tilde{X} = g(Y) = f^{-1}(Y) + r, \quad (2)$$

where  $r$  is the reconstruction residual. A true inverse  $f^{-1}(\cdot)$  does not generally exist, so SISR represents an ill-posed inverse problem.

Despite being ill-posed, the reason why SISR can successfully restore resolution is that both  $X$  and  $Y$  share information that can be represented in a low-dimensional manifold. A well-trained SISR model should be able to extract visual features from  $Y$  and map it into an image feature space. Then  $X$  can be reconstructed from the manifold with correct feature mapping. Dong et al. (2016a) have shown that CNNs have a built-in nature for the above processes. In a CNN-based SISR technique, all three different steps are trained together: feature extraction, manifold learning, and image reconstruction. This mingling of different components requires the network to extract the representative low-level feature, construct representative feature space, and precisely reconstruct images from features, which makes CNN based approach achieves state-of-art performance (Dong et al., 2016b; Kim et al., 2016).

### 3.2. GAN-based Super-Resolution

Most of the previous SISR approaches optimize the reconstruction by minimizing the voxel-wise difference ( $L_1$  or  $L_2$  loss) between  $\tilde{X}$  and  $X$ . However, Ledig et al. (2017) points out that merely taking care of local voxel-wise differences cast extreme difficulty in restoring important small details due to the ambiguity of the mapping between  $X$  and  $Y$ . We demonstrate one toy example in **Fig. 2**, where the HR image is  $2 \times 2$  down-sampled to an LR image, and the neighborhood is only in  $2 \times 2$  pixels. When only guided with  $L_1$  loss, the SR model doesn't have enough contextual information to recover local neighbors fully. By minimizing the Euclidean loss, it tends to average all possible HR candidates, resulting in a blurred output. However, if we put global perceptual constraints into the account, the SR model is guided by both local intensity information and patch-wise perceptual information, possibly making SR sharper and better-looking. However, such guiding is impossible to be handcrafted because there is no well-adapted mathematical definition of good perceptual quality for images. Based on this observation, Ledig et al. (2017) proposed to use a Generative Adversarial Network (GAN) for its unsupervised-learning potential of capturing perceptually important image features.




**Fig. 2.** An example when an SR model is optimized by  $L_1$  vs. perceptual loss in a  $2 \times 2$  neighborhood. The down-sampled LR is the same from two HR patches. Instead of voxel-wisely averaging all possible HR candidates which causes over-smoothing, GAN drives towards perceptual favorable SR solutions by taking account of other informations (i.e. positions) and features in the image manifolds.

The GAN framework proposed by Goodfellow et al. (2014) has two networks: a generator  $G$  and a discriminator  $D$ . The principle of a GAN is to train a  $G$  that generates fake images as real as possible, while simultaneously to train a  $D$  to distinguish the genuine of them. After training,  $D$  becomes very good at separating real and generated images, while the  $G$  learns to produce realistic-looking images by the "instruction" from  $D$ . GAN can model the image representation in an unsupervised manner that doesn't require a pre-designed objective. It is a perfect fit for a SISR. SRGAN (Ledig et al., 2017) was proposed and shows that the SR model yields unprecedented perceptual quality with the help of GAN.

However, training a GAN could be very challenging. The balance between  $G$  and  $D$  has to be carefully maintained so that both of them evolve together. Otherwise, if either side of the lever is too strong, the training quickly landslides to one side, resulting in an under-trained generator  $G$  (Salimans et al., 2016). A lot of efforts have been made to stabilize the GAN's training. However, those approaches are highly dependent on the specific network structure, and barely any research has investigated a 3D GAN network. Arjovsky et al. (2017) observed that the collapse of vanilla GAN training is caused by its optimization toward Kullback-Leibler (KL) divergence between the real and generated probability when there is little or no overlap between them, which is very common at the early stage of training; the gradient from  $D$  vanishes, which causes the training to halt. To address this issue, they proposed Wasserstein GAN (WGAN), whose objective is to minimize an efficient approximation of Earth Mover (EM) distance. They proved that this change could remove the difficult-to-achieved requirement for balancing  $D$  and  $G$ . The WGAN enables almost fail-free training in any situation while keeping the quality as good as a vanilla GAN. Additionally, the EM distance from  $D$  can also indicate the output image's quality, which is very useful for training.






**Fig. 3.** mDCSRN-GAN overview. The Generator is our proposed mDCSRN. The Discriminator is adapted from Ledig et al. (2017).

### 3.3. Proposed 3D Multi-Level Densely Connected Super-Resolution Network (mDCSRN)

Our proposed mDCSRN uses a DenseNet (Huang et al., 2017) as the starting point. By adding a multi-level densely connection and compressor in each Densely Connected Block (DenseBlock), our network is even more memory-efficient than the original DenseNet and provides excellent images in 3D SISR. An overview of our framework is shown in **Fig. 3**. All DenseUnits have a growth rate  $k = 12$ . We chose exponential linear units (ELU) (Clevert et al., 2015) as the activation layer to make use of negative values of normalized MRI. We placed a stem module that contains a convolution layer with  $2k$  filters before the feature mapping network, which is a set of densely connected DenseBlocks. The last part of our mDCSRN is the reconstruction module, which forms the final output. All convolutional layers are using  $3 \times 3 \times 3$  kernels, except those in the compressor within the DenseBlock and the direct combination layer in the reconstruction module, where kernel size is  $1 \times 1 \times 1$ . There is no up-sampling layer in mDCSRN. As the resolution loss in LR MRI is not in the spatial domain but the k-space, both LR and HR MRI are often generated with the same matrix size when directly fetched from a scanner. We want to discuss structure details as following:

**Fully Densely Connected Block.** The backbone of the mDCSRN is the DenseBlock from DenseNet (Huang et al., 2017). We fully connected all layers within DenseBlocks. It helps to increase feature sharing, making the neural network fewer parameters to keep the same representation capacity. As shown in **Fig. 4**, in our implementation, the input feature map is always directly connected to every convolutional layer, including the output within the DenseBlock, while in Tong et al. (2017) these connections are missing. Those direct links ensure that each DenseUnit can access not only preceding layers within the same




**Fig. 4.** Two connectivity ways of a DenseNet: (a) our proposed mDCSRN vs (b) SRDenseNet (Tong et al., 2017). Dense connections from the input (red lines in (a)) are missing in (b), which eliminates the direct link to the preceding DenseBlocks.

DenseBlock but also those in the preceding DenseBlocks, and lead to higher efficiency in parameter usage. To further reduce memory usage, as mentioned in DenseNet-bc (Huang et al., 2017), we also put a  $1 \times 1 \times 1$  bottleneck layer with 4k width before each  $3 \times 3 \times 3$  convolution when needed.

**Multiple Hierarchy Level with Fully Dense Connections.** Veit et al. (2016) found that Highway Network (Srivastava et al., 2015) and ResiNet (He et al., 2016) with skip connections act equally as an ensemble of multiple shallow networks with many paths instead of a giant deep network. Each small network processes some tasks on a different visual level depends on their position. This hierarchical structure harmonizes the animal’s visual system discovered by Hubel and Wiesel (1962), which might explain deep ResiNet’s excellent performance. As the links within a DenseNet are more effective than ResiNet, this effect is more obvious: all convolutional layer can access all other levels of information and contributes together to the final output. Hence, DenseNet SR is more powerful, as shown in SRDenseNet (Tong et al., 2017).

**Densely Connected DenseBlocks and Compressor.** Though a deep learning model with a single DenseBlock is already capable of providing high-quality SR images (Chen et al., 2018), a more sophisticatedly designed architecture still promises better performance. Yet even memory-efficient DenseNets have too many parameters when constructed in 3D. To reduce memory usage while keeping the inter-links strong, we followed the principles of DenseNet and proposed a multi-level densely connected structure. We grouped DenseUnits into DenseBlocks with extra levels of dense connections, as shown in **Fig. 3(G)**. Then a  $1 \times 1 \times 1$  convolutional layer (compressor) is applied before each DenseBlock with a fixed output filter number of 2k. According to (Szegedy et al., 2016), this compressor does not negatively affect performance but reduces the weights dramatically. We believe that it brings us at least two advantages: 1) It greatly lessens the parameter number and computation cost; 2) It evens out the weights of different DenseBlocks, forcing the model to focus on low-level



**Fig. 5.** Reconstruction network: (a) Directed Feature Combination as proposed in mDCSRN (b) Reconstruction with a bottleneck (8k) followed by a BatchNorm and convolutional layer as proposed in Tong et al. (2017).

features.

**Direct Feature Combination.** To further shrink down the model size and improve running speed, in the last module of mDCSRN, we replaced conventional spatial convolutional layers with a 1x1x1 convolutional layer to directly combine all feature maps to the final SR output. This reconstruction process acts as an adaptive feature selection to jointly fuse all the DenseBlock’s output. Besides efficiency, as a single DenseBlock is already powerful enough to produce high-quality SR images, our design boosts the ensemble effects of small networks dealing with different visual level information (Liu et al., 2016), which conceivably improves SR image quality.

**GAN-Guided Training (mDCSRN-GAN).** To achieve plausible-looking SR results, we utilized the adversarial loss from a discriminator in a GAN. The discriminator  $D$  is built based on the structures of the  $D$  in SRGAN (Ledig et al., 2017). For the type of GAN, we chose WGAN for its excellent stability. Moreover, we use the gradient penalty variant of WGAN, known as WGAN-GP (Gulrajani et al., 2017), to accelerate converging in training. As suggested by WGAN-GP, we replace the batch normalization(BN) layer with layer normalization(LN) in the discriminator  $D$ .

**Loss Function.** Our loss function is composed of two parts: intensity loss,  $loss_{int}$ , and adversarial loss from GAN’s discriminator,  $loss_{adv}$ :

$$loss = loss_{int} + \lambda loss_{adv}, \quad (3)$$

where  $\lambda$  is a hyper-parameter, set to 0.1 in experiments. We used the absolute different ( $L_1$  loss) between the network output SR and ground-truth HR as the intensity loss:

$$loss_{int} = loss_{L_1} = \frac{\sum_{z=1}^L \sum_{y=1}^H \sum_{x=1}^W |I_{x,y,z}^{HR} - I_{x,y,z}^{SR}|}{LHW}, \quad (4)$$

where  $I_{x,y,z}^{SR}$  is the SR and  $I_{x,y,z}^{HR}$  is the ground-truth image patch. WGAN’s discriminator loss is used as an additional loss in SRGAN network training:

$$loss_{adv} = loss_{WGAN,D} = -D_{WGAN,\theta}(I^{SR}), \quad (5)$$

where  $D_{WGAN,\theta}(I^{SR})$  is WGAN’s discriminator output digit for generated SR image patch.

### 3.4. LR Image Generation

An approach to generate LR images from original resolution HR images is required to evaluate the SISR technique. We follow the same steps as in Chen et al. (2018): 1) apply 3D FFT to transform HR image into k-space; 2) reduce the resolution by truncating outer part of k-space with a factor of 2x2 in both phase-encoding directions ( $2 \times 2 \times 1$  ratio in total); 3) convert back to image space by applying inverse FFT and then linearly interpolate to the original image size. This process mimics the actual acquisition of LR and HR images by MRI scanners.

## 4. Experiments

We first describe our experimental settings. Then we conduct a set of experiments to demonstrate that the proposed mDCSRN is not only memory-efficient but also provides state-of-the-art SR results by quantitative metrics. Next, we show that our mDCSRN-GAN provides encouraging qualitative results that are comparable with the ground-truth HR images, as demonstrated by the perceptual scores.

### 4.1. Settings

**Datasets.** To demonstrate the generalization of mDCSRN, we used the data from the Human Connectome Project (HCP) (Van Essen et al., 2013), which is a comprehensive publicly accessible brain MRI database with 1113 subjects. The 0.7 mm isotropic high-resolution 3D T1W images with a matrix size of  $320 \times 320 \times 256$  were acquired via Siemens 3T Prisma platform on multiple centers. The high-quality ground truth HR images with detailed small structures make this dataset a perfect case to test SISR approaches. The whole dataset is subject-wise split into 780 training, 111 validation, 111 evaluation, and 111 test samples. No subjects nor image patches are overlapped in any subsets. The validation set is used for monitoring and getting the best model checkpoint that has the highest performance during training, measured using mean square error (MSE) for non-GAN training, and EM-distance for GAN training. The evaluation set was used for hyper-parameter searching. The test set is only used for final performance analysis to avoid making model favorable to test data.

**Training Details.** The model was implemented in Tensorflow (Abadi et al., 2016) on a workstation with Nvidia GTX 1080 TI GPUs. For non-GAN networks, ADAM (Kingma and Ba) optimizer with a learning rate of  $10^{-4}$  was used to minimize the  $L_1$  loss. The batch size was set to 6. We followed a similar process of patching and data augmentation as in Chen et al. (2018), except, the patch size during training was set as  $40 \times 40 \times 40$ . We trained mDCSRN for 800k iterations, which is about 300 epochs, as 18 randomly sampled patches were fetched from a patient during training, lasting from 5 to 14 days depending on network size. For GAN experiments, we transfer the weights from well-trained mDCSRN above as the initial  $G$  of mDCSRN-GAN. We first trained  $D$  for the initial 10k steps without updating  $G$ . After then, for 5 iterations of training the

$D$ ,  $G$  was trained once. Additionally, after every 500 iterations of  $G$  training,  $D$  was trained for an extra 200 steps. It is solely to make sure  $D$  is always ahead of  $G$ , as suggested in WGAN (Arjovsky et al., 2017). Adam optimizer with  $5 \times 10^{-6}$  was used to optimize  $G$  for a total of 200k steps.

**SR Generation.** Once training was finished, LR images from the evaluation/test set were fed into the model to generate SR outputs. A patch size of  $70 \times 70 \times 70$  with a margin 3 was used in testing to avoid artifacts on the edges. The merging of the output patches was done without averaging. Because the batch size is 1 during testing, we set the batch normalization layers in the model to "train" mode instead of "test" mode for better estimation. We recorded the runtime speed on a single Nvidia GTX 1080 TI GPU.

**Quality Metrics.** To quantitatively measure mDCSRN’s recovery accuracy, we used three reference-based image similarity metrics: structural similarity index (SSIM) (Wang et al., 2004), peak signal to noise ratio (PSNR), and normalized root mean squared error (NRMSE). Numbers were calculated in the most resolution degraded cross-section ( $2 \times 2$ ) slice by slice. Scores were reported in its subject-wise slice-averaged numbers. For mDCSRN-GAN measurement, we list its numeric metrics as well. But we need to point out that PSNR could not fully represent the visual quality. Hence, we measured the perceptual quality via non-reference metrics: PIQE (Venkatanath et al., 2015), Ma’s score (Ma et al., 2017), NIQE (Mittal et al., 2012), and perceptual index (PI, used in PRIM-SR Challenge (Blau et al., 2018)). To efficiently calculate the perceptual scores, we only processed the 2D slices where the foreground (brain region) occupies more than 25% of the whole image. All perceptual scores were calculated in MATLAB R2019 software.

**Segmentation Evaluation.** In the testing stage, to further exemplify the benefits from our SR for the automatic medical image processing system, we conducted a fully automated segmentation on 159 brain tissues from a pre-trained high-performance neural network: HighRes3D (Li et al., 2017). We performed the test on the output of bicubic interpolation, SRResNet, mDCSRN *b8u4*, and mDCSRN-GAN *b8u4*. We first interpolated all images from the original  $0.7mm^3$  spatial resolution into  $1.0mm^3$  since the HighRes3D network was trained on the latter resolution. Then, we performed an N4 bias correction (Tustison et al., 2010) with ANTS (Avants et al., 2009) toolbox. Then we ran the inferences of HighRes3D on the NiftyNet (Gibson et al., 2018) open-platform. We used two similarity metrics, Dice Similarity Coefficient (DSC) (Sørensen, 1948) and Jaccard Index (JACC) (Jaccard, 1901), to quantitatively measure the agreement of segmentation between the up-sampled/super-resolution and the high-resolution images. Numbers were average among those 159 different anatomical structures.

#### 4.2. Results

We first demonstrate that the compressor in our multi-level densely connection does improve memory efficiency. We show that by replacing spatial convolutional layers with a single direct feature combination, we further reduce the model size without sacrificing performance. We show how the depth and width

Table 1: Ablation experiment results of mDCSRN on the **evaluation** set

	PSNR‡	SSIM‡	NRMSE†	#Parm	time(s)
Exp. 1			k=12		
<b>b1u16-r</b>	<b>35.84 ± 0.86</b>	0.9408 ± 0.0060	<b>0.0866 ± 0.0036</b>	0.35M	18.69
<b>b4u4-r</b>	<b>35.84 ± 0.86</b>	<b>0.9410 ± 0.0060</b>	<b>0.0866 ± 0.0037</b>	0.26M	<b>12.93</b>
<b>b1u12-r</b>	35.76 ± 0.86	0.9396 ± 0.0060	0.0874 ± 0.0036	0.25M	<b>12.74</b>
Exp. 2			k=12		
<b>b4u4-r</b>	35.84 ± 0.86	<b>0.9410 ± 0.0060</b>	0.0866 ± 0.0037	0.26M	12.93
<b>b4u4</b>	<b>35.88 ± 0.85</b>	0.9408 ± 0.0060	<b>0.0864 ± 0.0037</b>	<b>0.22M</b>	<b>12.06</b>
Exp. 3			k=12		
<b>b4u4</b>	35.88 ± 0.85	0.9408 ± 0.0060	0.0864 ± 0.0037	0.22M	<b>12.06</b>
<b>b6u4</b>	36.06 ± 0.86	0.9431 ± 0.0059	0.0845 ± 0.0037	0.35M	19.14
<b>b8u4</b>	<b>36.14 ± 0.87</b>	<b>0.9442 ± 0.0059</b>	<b>0.0836 ± 0.0037</b>	0.49M	28.52
Exp. 4			b4u4		
<b>k=8</b>	35.57 ± 0.85	0.9382 ± 0.0060	0.0894 ± 0.0036	<b>0.10M</b>	<b>7.95</b>
<b>k=12</b>	35.88 ± 0.85	0.9408 ± 0.0060	0.0864 ± 0.0037	0.22M	12.06
<b>k=16</b>	<b>35.96 ± 0.87</b>	<b>0.9424 ± 0.0059</b>	<b>0.0854 ± 0.0037</b>	0.41M	15.37
Exp. 5					
<b>b4u4k12</b>	<b>35.88 ± 0.85</b>	0.9408 ± 0.0060	0.0864 ± 0.0037	0.22M	<b>12.06</b>
<b>b8u4k8</b>	35.85 ± 0.86	<b>0.9415 ± 0.0059</b>	<b>0.0863 ± 0.0037</b>	0.22M	18.43
<b>b4u4k16</b>	35.96 ± 0.87	0.9424 ± 0.0059	0.0854 ± 0.0037	0.41M	<b>15.37</b>
<b>b8u4k12</b>	<b>36.14 ± 0.87</b>	<b>0.9442 ± 0.0059</b>	<b>0.0836 ± 0.0037</b>	0.49M	28.52

‡: The higher the better, †: The lower the better

**b**:# DenseBlock, **u**:# DenseUnit per Block, **k**: Growth rate **-r**: with reconstruction layer; default is using direct combination layer

of mDCSRN affect performance, and we compare mDCSRN with other popular SISR models. Qualitatively, we show the results from the mDCSRN-GAN side by side with other up-sampling methods. The mDCSRN-GAN provides realistic-looking images while running at the same time as our mDCSRN. We further investigate the perceptual quality with quantitative non-reference metrics. To demonstrate our model’s clinical value in automatic systems, we use the brain tissue segmentation as an example to demonstrate the benefits brought by SR models. Last, we show that in the real-world scan, our mDCSRN-GAN exhibits its fantastic stability across different platforms.

**Multi-Level Connectivity and Compressor.** As shown in **Table 1** Exp. 1, with the same total number of DenseUnit, mDCSRN b4u4-r had fewer parameters, ran faster, and achieved the same performance as the original DenseNet design b1u16-r; with the same amount of parameters, b4u4-r significantly outperformed b1u12-r; proving that multi-level connectivity and compressor together helped improve memory efficiency and runtime speed.

**Direct Feature Combination vs Extra Reconstruction Layer.** As shown in **Table 1** Exp. 2, with the same depth, b4u4 with our introduced direct feature combination achieved similar to slightly better performance than b4u4-r with reconstruction layers while decreasing model size by 15%.

**Depth vs Width.** The results with different depth and width configuration

Table 2: mDCSRN vs. interpolation and previous CNN based SISRs on the **test** set

Intensity-based similarity metrics					
	PSNR‡	SSIM‡	NRMSE‡	#parm	time(s)
<b>NN</b>	29.48 ± 0.81	0.8219 ± 0.0113	0.2007 ± 0.0071	N/A	N/A
<b>Bicubic</b>	30.30 ± 0.82	0.8420 ± 0.0105	0.1830 ± 0.0067	N/A	N/A
<b>FSRCNN</b>	34.33 ± 0.81	0.9207 ± 0.0062	0.1142 ± 0.0050	<b>0.06M</b>	15.57
<b>SRResNet</b>	36.09 ± 0.82	0.9425 ± 0.0052	0.0939 ± 0.0043	2.01M	107.16
<b>SRDenseNet</b>	35.93 ± 0.82	0.9413 ± 0.0052	0.0955 ± 0.0044	0.39M	17.95
<b>b4u4k12</b>	36.08 ± 0.82	0.9418 ± 0.0052	0.0935 ± 0.0044	0.23M	<b>12.54</b>
<b>b6u4k12</b>	36.31 ± 0.82	0.9438 ± 0.0051	0.0915 ± 0.0043	0.35M	19.60
<b>b8u4k12</b>	<b>36.39 ± 0.82</b>	<b>0.9448 ± 0.0050</b>	<b>0.0906 ± 0.0043</b>	0.49M	27.86


  

Perceptual quality metrics				
	PIQE‡	NIQE‡	MA’s Score‡	PI‡
<b>Bicubic</b>	99.54 ± 1.39	5.53 ± 0.13	3.20 ± 0.06	6.17 ± 0.08
<b>SRResNet</b>	80.85 ± 7.51	5.92 ± 0.19	<b>5.06 ± 0.03</b>	5.43 ± 0.08
<b>mDCSRN <i>b8u4</i></b>	81.04 ± 7.79	6.01 ± 0.19	<b>5.06 ± 0.03</b>	5.48 ± 0.09
<b>mDCSRN-GAN <i>b8u4</i></b>	<b>71.86 ± 7.24</b>	<b>5.00 ± 0.16</b>	5.04 ± 0.04	<b>4.98 ± 0.07</b>
<b>Original Resolution</b>	<b>69.70 ± 7.02</b>	5.52 ± 0.14	5.05 ± 0.03	5.23 ± 0.06

‡: The higher the better, †: The lower the better

are shown in **Table 1** Exp. 3 and Exp. 4. The performance was improved by either making the network deeper or wider, at the cost of more extensive memory consumption and slower inference speed. As shown in **Table 1** Exp. 5, when models are in a similar size, the deeper network, the better the performance. Although the weight-saving mechanism is more effective in the deep and narrow network, it runs slower, due to the extra computational cost from additional bottleneck layers. Therefore, given a fixed memory constraint, a shallow mDCSRN is preferable for a fast application, while a deep mDCSRN is excellent for better results.

**Baseline.** As baseline models, FSRCNN (Dong et al., 2016b), SRResNet (Ledig et al., 2017), and SRDenseNet (Tong et al., 2017) were implemented and extended to 3D. As there is no image-size changing in our SISR, the up-sampling CNNs (transposed-convolutional layers or sub-pixel layers) in those original designs were replaced with the same scale convolutional layers. For SRDenseNet, we adjusted the hyperparameters as similar as possible to mDCSRN *b8u4* (i.e. reduced DenseUnit number from 8 to 4, changed activation function to ELU, and set growth-rate  $k=12$ ). All models were trained for 300 epochs. With respect to quantitative similarity metrics, as shown in **Table 2**, the lightest mDCSRN *b4u4* ran fastest among all CNN approaches with competitive results. The deepest mDCSRN *b8u4* as shown in **Fig. 6** outperformed all previous SISR approaches by a considerable margin. It did run slower than SRDenseNet but was still 4x faster than the SRResNet. Both SRDenseNet and mDCSRN *b6u4* are similar in model size and running speed, but the later significantly outperformed the former, proving the advantage of our efficient architecture design.



**Fig. 6.** Example results from the test set of Nearest Neighbor, SRResNet, mDCSRN *b8u4*, mDCSRN-GAN *b8u4* in the  $2 \times 2$  resolution degraded plane. PSNR and SSIM of this subject are shown on the top. Despite performing worse in PSNR and SSIM, GAN SR images appear to have recovered more spatial details.

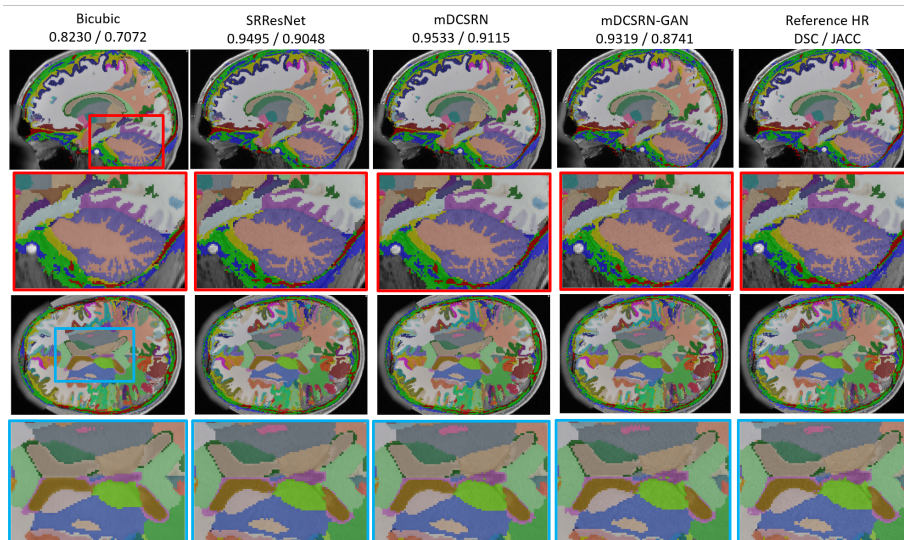
Table 3: Segmentation accuracy on the **test** set

	Bicubic	SRResNet	mDCSRN <i>b8u4</i>	mDCSRN-GAN <i>b8u4</i>
<b>DSC</b>	$0.810 \pm 0.107$	$0.946 \pm 0.047$	$0.949 \pm 0.045$	$0.929 \pm 0.056$
<b>JACC</b>	$0.691 \pm 0.127$	$0.899 \pm 0.061$	$0.905 \pm 0.059$	$0.871 \pm 0.071$

**Perceptual Quality.** An example output is shown in **Fig. 6**. mDCSRN *b8u4* provides slightly better SR reconstruction accuracy than SRResNet, but it is mDCSRN-GAN *b8u4* that more closely shapes the small vessel pointed by the red arrows. Though mDCSRN-GAN’s PSNR is lower than its non-GAN sibling, it provides more structural details that are more plausible by the human eye. As shown in **Table 2**, the quantitative perceptual quality numbers suggest that while non-GAN SR shows slightly closer to HR only in the MA’s metric, the GAN SR model shows much better performance in all other three measurements. GAN even obtained a higher score in NIQE and PI than HR, since SR images were generated from less noisy LR input, making the SR more plausible for noise-sensitive perceptual metrics. Wang et al. (2018) has shown similar results in their SR and HR perceptual comparison.

**Segmentation Task.** We investigated the segmentation results on the output of interpolation, SRResNet, mDCSRN, and mDCSRN-GAN. As shown in the **Table 3**, the segmentation results are more aligned with similarity metrics. That’s because the segmentation task is more focused on the contrast instead of realistic patterns. Overall, segmentation from the SR models’ output is more consistent with the segmentation of the original resolution. The high overlapping between those two indicates that segmentation on SR images are not be





**Fig. 7.** An sample test case of segmentation from HighRes3DNet (Li et al., 2017) on the output of bicubic interpolation, SRResNet, mDCSRN *b8u4*, mDCSRN-GAN *b8u4*, and Original Resolution. Average similarity metrics of this subject among 159 structures are shown on the top.


Table 4: Perceptual image quality metrics in **real-world** scans (N=7)

	NIQE <sup>†</sup>	MA’s Score <sup>‡</sup>	PI <sup>†</sup>
<b>Low-resolution</b>	$7.44 \pm 0.52$	$3.83 \pm 0.19$	$6.80 \pm 0.18$
<b>mDCSRN-GAN</b>	$5.66 \pm 0.77$	$4.92 \pm 0.05$	$5.37 \pm 0.39$
<b>Full-resolution</b>	$5.17 \pm 0.16$	$4.97 \pm 0.03$	$5.10 \pm 0.07$

<sup>‡</sup>: The higher the better, <sup>†</sup>: The lower the better

greatly different from those on HR. An example is shown in **Fig. 7**.

**Prospective MR Scans.** Additionally, we also performed a real-world test on seven volunteers in our on-site 3T Siemens Verio MRI scanner, which is different to those Prisma scanners utilized in the HCP dataset. We followed the same protocol as in Van Essen et al. (2013) except for reducing the phase encoding and slice resolution by half, which effectively reduced spatial resolution by 4x. As shown in **Fig. 8** and **Table 4**, the mDCSRN-GAN model showed excellent ability in recovering edge details that hardly seen in the fast low-resolution scan. Besides noticeable sharpness improvement, the SR output seems to have a lower noise level and cleaner image than the original full-resolution scan because of the low-resolution image that has a better SNR than HR. It’s an extra gain from super-resolution techniques in addition to the time- and cost-saving. As the real scan was performed on a completely different machine on a different site and subject, the noise pattern and image quality were considerably different than the training dataset. It displays our model’s robustness and performance



**Fig. 8.** Two **real-world** examples are shown in  $2 \times 2$  resolution-reduced plane. There are slight mismatches between LR and HR, because they are from two separate scans. These scans were done on a different version of Siemens MRI scanner at Cedars-Sinai Medical Center. mDCSRN-GAN provides a comparable image quality to high-resolution scan.

in a real-world scenario.

## 5. Conclusions

In this paper, we developed and evaluated a highly efficient architecture mDCSRN for 3D MRI SISR. We showed that the proposed mDCSRN could outperform common existing methods in voxel-based similarity metrics and segmentation accuracy with a smaller model size. We also demonstrated that with GAN-guided training, our mDCSRN-GAN could successfully recover fine details and further improve perceptual quality. Testing on prospectively acquired data showed that our model is capable of real-world clinical application. In summary, the new technique would allow a 4-fold reduction in scan time with minimal loss in image details and perceptual quality, which would substantially improve the clinical practicality of high-resolution MRI.

## References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al., 2016. TensorFlow: A system for large-scale machine learning., in: OSDI, pp. 265–283.
- Arjovsky, M., Chintala, S., Bottou, L., 2017. Wasserstein generative adversarial networks, in: International Conference on Machine Learning, pp. 214–223.
- Avants, B.B., Tustison, N., Song, G., 2009. Advanced normalization tools (ANTS). *Insight j* 2, 1–35.
- Blau, Y., Mechrez, R., Timofte, R., Michaeli, T., Zelnik-Manor, L., 2018. The 2018 PIRM challenge on perceptual image super-resolution, in: Proceedings of the European Conference on Computer Vision (ECCV), pp. 0–0.
- Chen, Y., Xie, Y., Zhou, Z., Shi, F., Christodoulou, A.G., Li, D., 2018. Brain MRI super resolution using 3D deep densely connected neural networks, in: 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018), IEEE. pp. 739–742.
- Clevert, D.A., Unterthiner, T., Hochreiter, S., 2015. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289* .
- Dong, C., Loy, C.C., He, K., Tang, X., 2016a. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence* 38, 295–307.
- Dong, C., Loy, C.C., Tang, X., 2016b. Accelerating the super-resolution convolutional neural network, in: European Conference on Computer Vision, Springer. pp. 391–407.
- Gibson, E., Li, W., Sudre, C., Fidon, L., Shakir, D.I., Wang, G., Eaton-Rosen, Z., Gray, R., Doel, T., Hu, Y., et al., 2018. NiftyNet: a deep-learning platform for medical imaging. *Computer methods and programs in biomedicine* 158, 113–122.
- Glasner, D., Bagon, S., Irani, M., 2009. Super-resolution from a single image, in: Computer Vision, 2009 IEEE 12th International Conference on, IEEE. pp. 349–356.
- Glorot, X., Bengio, Y., 2010. Understanding the difficulty of training deep feedforward neural networks, in: Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, pp. 249–256.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y., 2014. Generative adversarial nets, in: Advances in neural information processing systems, pp. 2672–2680.

- Greenspan, H., 2008. Super-resolution in medical imaging. *The Computer Journal* 52, 43–63.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.C., 2017. Improved training of wasserstein gans, in: *Advances in Neural Information Processing Systems*, pp. 5769–5779.
- 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.
- Huang, G., Liu, Z., Weinberger, K.Q., van der Maaten, L., 2017. Densely connected convolutional networks, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, p. 3.
- Hubel, D.H., Wiesel, T.N., 1962. Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *The Journal of physiology* 160, 106–154.
- Ioffe, S., Szegedy, C., 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift, in: *International conference on machine learning*, pp. 448–456.
- Jaccard, P., 1901. Étude comparative de la distribution florale dans une portion des alpes et des jura. *Bull Soc Vaudoise Sci Nat* 37, 547–579.
- Johnson, J., Alahi, A., Fei-Fei, L., 2016. Perceptual losses for real-time style transfer and super-resolution, in: *European Conference on Computer Vision*, Springer. pp. 694–711.
- Kim, J., Kwon Lee, J., Mu Lee, K., 2016. Accurate image super-resolution using very deep convolutional networks, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1646–1654.
- Kingma, D.P., Ba, J., . Adam: A method for stochastic optimization, in: *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*, arXiv preprint arXiv.
- Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., et al., 2017. Photo-realistic single image super-resolution using a generative adversarial network, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4681–4690.
- Li, W., Wang, G., Fidon, L., Ourselin, S., Cardoso, M.J., Vercauteren, T., 2017. On the compactness, efficiency, and representation of 3D convolutional networks: brain parcellation as a pretext task, in: *International Conference on Information Processing in Medical Imaging*, Springer. pp. 348–360.

- Lim, B., Son, S., Kim, H., Nah, S., Lee, K.M., 2017. Enhanced deep residual networks for single image super-resolution, in: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, p. 3.
- Litjens, G., Kooi, T., Bejnordi, B.E., Setio, A.A.A., Ciompi, F., Ghafoorian, M., van der Laak, J.A., van Ginneken, B., Sánchez, C.I., 2017. A survey on deep learning in medical image analysis. *Medical image analysis* 42, 60–88.
- Liu, D., Wang, Z., Nasrabadi, N., Huang, T., 2016. Learning a mixture of deep networks for single image super-resolution, in: *Asian Conference on Computer Vision, Springer*. pp. 145–156.
- Ma, C., Yang, C.Y., Yang, X., Yang, M.H., 2017. Learning a no-reference quality metric for single-image super-resolution. *Computer Vision and Image Understanding* 158, 1–16.
- Mittal, A., Soundararajan, R., Bovik, A.C., 2012. Making a completely blind image quality analyzer. *IEEE Signal Processing Letters* 20, 209–212.
- Oktay, O., Bai, W., Lee, M., Guerrero, R., Kamnitsas, K., Caballero, J., de Marvao, A., Cook, S., ORegan, D., Rueckert, D., 2016. Multi-input cardiac image super-resolution using convolutional neural networks, in: *International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer*. pp. 246–254.
- Pang, J., Chen, Y., Fan, Z., Nguyen, C., Yang, Q., Xie, Y., Li, D., 2016. High efficiency coronary MR angiography with nonrigid cardiac motion correction. *Magnetic Resonance in Medicine* 76, 1345–1353.
- Park, S.C., Park, M.K., Kang, M.G., 2003. Super-resolution image reconstruction: a technical overview. *IEEE signal processing magazine* 20, 21–36.
- Pham, C.H., Ducournau, A., Fablet, R., Rousseau, F., 2017. Brain MRI super-resolution using deep 3D convolutional networks, in: *Biomedical Imaging (ISBI 2017), 2017 IEEE 14th International Symposium on, IEEE*. pp. 197–200.
- Pruessner, J.C., Li, L.M., Serles, W., Pruessner, M., Collins, D.L., Kabani, N., Lupien, S., Evans, A.C., 2000. Volumetry of hippocampus and amygdala with high-resolution MRI and three-dimensional analysis software: minimizing the discrepancies between laboratories. *Cerebral Cortex* 10, 433–442.
- Rudin, L.I., Osher, S., Fatemi, E., 1992. Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena* 60, 259–268.
- Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X., 2016. Improved techniques for training gans, in: *Advances in Neural Information Processing Systems*, pp. 2234–2242.

- Shi, F., Cheng, J., Wang, L., Yap, P.T., Shen, D., 2015. LRTV: MR image super-resolution with low-rank and total variation regularizations. *IEEE transactions on medical imaging* 34, 2459–2466.
- Sørensen, T., 1948. A method of establishing groups of equal amplitude in plant sociology based on similarity of species content and its application to analyses of the vegetation on danish commons. *Biologiske Skrifter* 5, 1–34.
- Srivastava, R.K., Greff, K., Schmidhuber, J., 2015. Training very deep networks, in: *Advances in neural information processing systems*, pp. 2377–2385.
- Stucht, D., Danishad, K.A., Schulze, P., Godenschweger, F., Zaitsev, M., Speck, O., 2015. Highest resolution in vivo human brain MRI using prospective motion correction. *PloS one* 10, e0133921.
- Sun, S., Chen, W., Wang, L., Liu, X., Liu, T.Y., 2016. On the depth of deep neural networks: A theoretical view., in: *AAAI*, pp. 2066–2072.
- Szegedy, C., Ioffe, S., Vanhoucke, V., Alemi, A.A., 2017. Inception-v4, inception-resnet and the impact of residual connections on learning., in: *AAAI*, p. 12.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z., 2016. Rethinking the inception architecture for computer vision, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2818–2826.
- Tai, Y., Yang, J., Liu, X., 2017. Image super-resolution via deep recursive residual network, in: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Tong, T., Li, G., Liu, X., Gao, Q., 2017. Image super-resolution using dense skip connections, in: *2017 IEEE International Conference on Computer Vision (ICCV)*, IEEE. pp. 4809–4817.
- Tustison, N.J., Avants, B.B., Cook, P.A., Zheng, Y., Egan, A., Yushkevich, P.A., Gee, J.C., 2010. N4ITK: improved N3 bias correction. *IEEE transactions on medical imaging* 29, 1310.
- Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K., Consortium, W.M.H., et al., 2013. The WU-Minn human connectome project: an overview. *Neuroimage* 80, 62–79.
- Veit, A., Wilber, M.J., Belongie, S., 2016. Residual networks behave like ensembles of relatively shallow networks, in: *Advances in Neural Information Processing Systems*, pp. 550–558.
- Venkatanath, N., Praneeth, D., Bh, M.C., Channappayya, S.S., Medasani, S.S., 2015. Blind image quality evaluation using perception based features, in: *2015 Twenty First National Conference on Communications (NCC)*, IEEE. pp. 1–6.

- Wang, S., Su, Z., Ying, L., Peng, X., Zhu, S., Liang, F., Feng, D., Liang, D., 2016. Accelerating magnetic resonance imaging via deep learning, in: Biomedical Imaging (ISBI), 2016 IEEE 13th International Symposium on, IEEE. pp. 514–517.
- Wang, X., Yu, K., Wu, S., Gu, J., Liu, Y., Dong, C., Qiao, Y., Change Loy, C., 2018. ESRGAN: Enhanced super-resolution generative adversarial networks, in: Proceedings of the European Conference on Computer Vision (ECCV), pp. 0–0.
- Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P., 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* 13, 600–612.
- Xie, L., Wisse, L.E., Das, S.R., Wang, H., Wolk, D.A., Manjón, J.V., Yushkevich, P.A., 2016. Accounting for the confound of meninges in segmenting entorhinal and perirhinal cortices in T1-weighted MRI, in: International Conference on Medical Image Computing and Computer-assisted Intervention, Springer. pp. 564–571.
- Yang, C.Y., Ma, C., Yang, M.H., 2014. Single-image super-resolution: A benchmark, in: European Conference on Computer Vision, Springer. pp. 372–386.
- Yang, H.J., Sharif, B., Pang, J., Kali, A., Bi, X., Cokic, I., Li, D., Dharmakumar, R., 2016. Free-breathing, motion-corrected, highly efficient whole heart T2 mapping at 3T with hybrid radial-cartesian trajectory. *Magnetic Resonance in Medicine* 75, 126–136.
- You, C., Li, G., Zhang, Y., Zhang, X., Shan, H., Li, M., Ju, S., Zhao, Z., Zhang, Z., Cong, W., et al., 2019. CT super-resolution GAN constrained by the identical, residual, and cycle learning ensemble (GAN-CIRCLE). *IEEE Transactions on Medical Imaging* 39, 188–203.
- Zhou, Z., Nguyen, C., Chen, Y., Shaw, J.L., Deng, Z., Xie, Y., Dawkins, J., Marbán, E., Li, D., 2017. Optimized cest cardiovascular magnetic resonance for assessment of metabolic activity in the heart. *Journal of Cardiovascular Magnetic Resonance* 19, 95.