



# HHS Public Access

Author manuscript

*Med Image Comput Comput Assist Interv.* Author manuscript; available in PMC 2021 October 27.

Published in final edited form as:

*Med Image Comput Comput Assist Interv.* 2021 ; 12906: 430–440. doi:10.1007/978-3-030-87231-1\_42.

## MRI Super-Resolution Through Generative Degradation Learning

Yao Sui<sup>1,2</sup>, Onur Afacan<sup>1,2</sup>, Ali Gholipour<sup>1,2</sup>, Simon K. Warfield<sup>1,2</sup>

<sup>1</sup>Harvard Medical School, Boston, MA, USA

<sup>2</sup>Boston Children's Hospital, Boston, MA, USA

### Abstract

Spatial resolution plays a critically important role in MRI for the precise delineation of the imaged tissues. Unfortunately, acquisitions with high spatial resolution require increased imaging time, which increases the potential of subject motion, and suffers from reduced signal-to-noise ratio (SNR). Super-resolution reconstruction (SRR) has recently emerged as a technique that allows for a trade-off between high spatial resolution, high SNR, and short scan duration. Deconvolution-based SRR has recently received significant interest due to the convenience of using the image space. The most critical factor to succeed in deconvolution is the accuracy of the estimated blur kernels that characterize how the image was degraded in the acquisition process. Current methods use handcrafted filters, such as Gaussian filters, to approximate the blur kernels, and have achieved promising SRR results. As the image degradation is complex and varies with different sequences and scanners, handcrafted filters, unfortunately, do not necessarily ensure the success of the deconvolution. We sought to develop a technique that enables accurately estimating blur kernels from the image data itself. We designed a deep architecture that utilizes an adversarial scheme with a generative neural network against its degradation counterparts. This design allows for the SRR tailored to an individual subject, as the training requires the scan-specific data only, i.e., it does not require auxiliary datasets of high-quality images, which are practically challenging to obtain. With this technique, we achieved high-quality brain MRI at an isotropic resolution of 0.125 cubic mm with six minutes of imaging time. Extensive experiments on both simulated low-resolution data and clinical data acquired from ten pediatric patients demonstrated that our approach achieved superior SRR results as compared to state-of-the-art deconvolution-based methods, while in parallel, at substantially reduced imaging time in comparison to direct high-resolution acquisitions.

### Keywords

MRI; Super-resolution; Deep learning

---

yao.sui@childrens.harvard.edu .

**Electronic supplementary material** The online version of this chapter ([https://doi.org/10.1007/978-3-030-87231-1\\_42](https://doi.org/10.1007/978-3-030-87231-1_42)) contains supplementary material, which is available to authorized users.

## 1 Introduction

MRI is critically important in clinical and scientific research studies. High spatial resolution in MRI allows for the precise delineation of the imaged tissues. However, the high spatial resolution requires a long scan time, and in turn unfortunately increases the potential of subject motion [1, 22] and reduces the signal-to-noise ratio [2, 16]. Super-resolution reconstruction (SRR) has recently emerged as a technique of post-acquisition processing, which allows for obtaining MRI images at a high spatial resolution, high SNR, and short scan time [7–9, 19, 26]. Among these SRR techniques, deconvolution-based methods [16, 20, 21, 23, 24] have recently received significant interest due to the convenience of using the image space only. However, the deconvolution requires an accurate estimate in the blur kernel that characterizes the degradation process of the acquired images, which are usually difficult to obtain. Current methods take handcrafted filters as an approximation of the blur kernel, such as Gaussian filters [7, 20]. As the image degradation is complex and varies with different sequences and scanners, however, handcrafted approximations do not necessarily ensure the success of the deconvolution, and even possibly lead to unusable reconstructions.

Blur kernel estimation is often used in natural image deblurring, also known as blind deconvolution [3], in particular equipped with deep neural networks [6]. Unfortunately, no techniques focus on blur kernel estimation for deconvolution-based SRR in MRI. Although deep learning-based techniques have been widely used in natural image super-resolution, the majority of deep SRR methods for MRI are performed with 2D slices [4, 28, 30] as the large-scale, auxiliary datasets of high-quality volumetric MRI images are practically challenging to obtain at a high resolution and suitable SNR. It has been shown that the deep SRR model learned on volumetric data achieved better results than on 2D slice data [15].

We sought to develop a methodology that allows for accurately estimating blur kernels from the image data itself. We designed a deep architecture that utilizes an adversarial scheme with a generative neural network against its degradation counterparts. Our design enables the SRR tailored to an individual subject, as the training of our deep SRR model requires the scan-specific image data only. We achieved high-quality brain MRI images through our SRR method at an isotropic resolution of 0.125 cubic mm with six minutes of imaging time. We assessed our method on simulated low-resolution (LR) image data and applied it to ten pediatric patients. Experimental results show that our approach achieved superior SRR as compared to state-of-the-art deconvolution-based methods, while in parallel, at substantially reduced imaging time in comparison to direct high-resolution (HR) acquisitions.

## 2 Methods

### 2.1 Theory

It is difficult for SRR to enhance 2D in-plane or true 3D MRI resolution due to the frequency encoding scheme [14, 18], but effective to improve through-plane resolution for 2D slice stacks [10]. Therefore, we focus on reducing the slice thickness of LR scans<sup>1</sup> with large matrix size and thick slices. The thick slices lead to reduced scan time and increased

---

<sup>1</sup>We refer to a 2D slice stack as a scan or an image hereafter.

SNR. Thus, we can acquire a set of LR images with different orientations to facilitate the SRR to capture the HR signals distributed in different directions, while keeping the total scan time short.

**Forward Model.**—Let  $x$  denote the HR reconstruction and  $\{y_j\}_{j=1}^n$  denote the  $n$  acquired LR images. The forward model that describes the acquisition process is formulated as

$$y_j = D_j B_j T_j x + v, \quad j = 1, 2, \dots, n. \quad (1)$$

The transform  $T_j$  compensates for subject motion and is obtained by aligning the  $n$  LR images together.  $B_j$  describes the image degradation defined by a blur kernel.  $D_j$  is the downsampling operator that discards a certain amount of data from the HR signal. The noise  $v$  is additive and Gaussian when  $\text{SNR} > 3$  [11].

**Generative Degradation Learning.**—Our SRR targets a joint estimate in both the HR reconstruction and blur kernel. We use a generative neural network to generate an HR estimate by  $x = f_\theta(z_x)$ , and  $n$  generative networks to generate  $n$  blur kernels by  $B_j = g_{\omega_j}(z_{b_j})$ . These generations are constrained by the image degradation process as described in the forward model. Therefore, the generative degradation learning is found by

$$\begin{aligned} \min_{x, \theta, B_j, \omega_j} \quad & \ell(x - f_\theta(z_x)) + \sum_{j=1}^n \ell_{b_j}(B_j - g_{\omega_j}(z_{b_j})), \\ \text{s.t.} \quad & y_j = D_j B_j T_j x + v, \quad j = 1, 2, \dots, n. \end{aligned} \quad (2)$$

where  $f_\theta$  and  $g_{\omega_j}$  are nonlinear functions that generate data, parameterized by  $\theta$  and  $\omega_j$ , respectively,  $\ell$  and  $\ell_{b_j}$  are the loss functions in the optimization, and  $z_x$  and  $z_{b_j}$  are the initial guesses for  $x$  and  $B_j$ , respectively. The function  $f_\theta$  here is also known as the deep image prior [25]. When the loss functions are  $\ell_2$  loss, we can substitute  $x$  and  $B_j$  with the generative functions, and the optimization is then re-formulated as

$$\min_{\theta, \omega_j} \sum_{j=1}^n \left\| y_j - D_j g_{\omega_j}(z_{b_j}) T_j f_\theta(z_x) \right\|_2^2 + \lambda \mathcal{R}_{TV}(f_\theta(z_x)), \quad (3)$$

with a total variation (TV) [17] term imposed to regularize the HR reconstruction for edge sharpness, and a weight parameter  $\lambda > 0$ .

We implement these generative functions by deep neural networks.  $f_\theta$  is realized by a 3D UNet-like encoder-decoder network [31].  $g_{\omega_j}$  is implemented by a fully connected network containing four hidden linear layers and one Tanh layers. Each layer is followed by a dropout regularization and a ReLU activation. The architecture of our generative degradation networks (GDN) is shown in Fig. 1. A degradation network comprises a generative function for the blur kernel in combination with the constraint delivered by the forward model. We solve the optimization by an Adam algorithm [12]. The training for the GDN is an optimization problem. It is thus carried out on the scan-specific LR images  $y_j$  themselves only, and in turn, allows for the SRR tailored to an individual subject.

The HR reconstruction is obtained directly from  $x = f_{\theta}(z_x)$  once the GDN model has been trained. Also, to ensure the appropriate scale of the voxel intensity, a standard TV-based SRR with the learned blur kernels  $B_j$  can be optionally applied to obtain the HR reconstruction.

## 2.2 GDN-Based SRR

Since we focus on enhancing the through-plane resolution, the degradation process is assumed to be associated with the blurs from the slice excitation and downsampling. The downsampling is carried out by truncating the high frequency in the Fourier domain, and it is thus determined by a sinc low-pass filter. The slice excitation is characterized by the slice profile that is generated by the radio frequency (RF) pulse during scans. Therefore, the blur kernel we estimate is the convolution between the slice profile and the sinc filter.

The motion compensation is implemented by a rigid body image registration as we focus on brain MRI in this work. We interpolate the LR images to those with the resolution and size the same as the HR reconstruction, and align them together. The obtained transformations are used as  $T_j$  in the forward model.

We initialize the blur kernels as a sinc filter convolved with a Gaussian low-pass filter whose full width at half maximum (FWHM) is equal to the slice thickness of the LR image over that of the HR reconstruction. We compute an initial guess of the HR reconstruction by a standard TV-based SRR [16] with the above initialized blur kernels. We set  $\lambda = 0.01$  in Eq. (3) according to our empirical results. We run 4k iterations with a learning rate of 0.01 to minimize Eq. (3). It took about 2.5 h to reconstruct an HR image of size  $384 \times 384 \times 384$  voxels from three LR scans on an NVIDIA Titan RTX GPU with PyTorch [13].

## 2.3 Materials

We simulated an LR dataset based on the Human Connectome Project (HCP) database [5]. We randomly selected forty subjects from the HCP database, including five T1w and five T2w HR images with an isotropic resolution of 0.7 mm as the ground truths. We generated four blur kernels, as shown in Figs. 2(b) and 2(c), based on a Dirac pulse, a Gaussian, a sinc, and a boxcar functions depicted in Fig. 2. We simulated four groups of LR images according to the four types of blur kernels. We downsampled each ground truth image to three LR images in the complementary planes with a slice thickness of 2 mm after convolving it with the blur kernel in its group. Gaussian noise was added in each LR image with a standard deviation of 5% of maximum voxel intensity.

We acquired thirty LR T2-FSE scans from ten pediatric patients - three images in complementary planes per patient - on a 3T scanner. The field of view (FOV) was  $192 \text{ mm} \times 192 \text{ mm}$ , matrix size was  $384 \times 384$ , slice thickness was 2 mm, TE/TR = 93/12890 ms with an echo train length of 16 and a flip angle of  $160^\circ$ . It took about two minutes in acquiring an LR image.

## 2.4 Experimental Design

As our SRR is an unsupervised deconvolution-based approach, we compared our approach to state-of-the-art methods in the same category, including a TV-based method [16], a gradient guidance regularized (GGR) method [20], and a TV-based method with joint blur kernel learning (TV-BKL) implemented by ourselves. The parameters of the baselines were set according to their best PSNR.

**Experiment 1: Assessments on Simulated Data.**—The goal of this experiment is two-fold: to assess the accuracy of the blur kernel estimates and to demonstrate that the estimated blur kernels lead to superior HR reconstructions. We reconstructed each HR image at an isotropic resolution of 0.7 mm from three LR images on the HCP dataset. We investigated the estimation error defined by the  $\ell_2$  norm of the difference between the estimated kernel and ground truth over the  $\ell_2$  norm of ground truth  $\|k_{est} - k_{gt}\|_2 / \|k_{gt}\|_2$ . To evaluate the reconstruction quality, we investigated the PSNR and SSIM [27] of the HR reconstructions.

**Experiment 2: Assessments on Clinical Data.**—This experiment aimed at assessing the applicability of our SRR for clinical use. We reconstructed each HR image at an isotropic resolution of 0.5 mm on the clinical dataset. We evaluated the sharpness of these reconstructions by average edge strength [29]. We checked the estimated blur kernels and qualitatively assessed these reconstructions.

**Experiment 3: Impact on Deconvolution-Based Algorithms.**—This experiment aimed at evaluating the impact of our generative degradation learning on deconvolution-based algorithms. We expected our estimated blur kernels can improve the TV and GGR algorithms in their SRR quality on the HCP dataset.

## 3 Results

### Experiment 1: Assessments on Simulated Data.

Figure 3 shows the estimation errors in the blur kernels obtained from our method and TV-BKL. Our method, SR-GDN, considerably outperformed TV-BKL on the HCP dataset. The average errors of SR-GDN and TV-BKL were respectively  $10.1\% \pm 4.6\%$  and  $25.8\% \pm 12.0\%$ . Two-sample  $t$ -test showed that SR-GDN offered significantly lower errors on the simulated data than TV-BKL at the 5% significance level ( $p = 2.16e^{-118}$ ).

Figure 4 shows the spectrum of true and estimated blur kernels from SR-GDN and TV-BKL on four representative simulations. Our estimates (SR-GDN) closely followed the true kernels and offered higher accuracy than TV-BKL with all types of kernels on the HCP dataset.

Figure 5 shows the quantitative assessment of our approach (SR-GDN) and the three baselines on the HCP dataset in terms of PSNR and SSIM. Our method achieved an average PSNR at  $41.2 \text{ dB} \pm 3.74 \text{ dB}$  and SSIM at  $0.98 \pm 0.01$ , which were considerably higher than obtained from the three base-lines. Two-sample  $t$ -tests showed that SR-GDN significantly

outperformed the three baselines in terms of SSIM at the 5% significance level with  $p = 2.36e^{-18}$  against TV,  $4.37e^{-18}$  against GGR, and  $1.81e^{-7}$  against TV-BKL.

### Experiment 2: Assessments on Clinical Data.

Figures 6(a) and 6(b) show the mean and standard deviation of the blur kernels estimated by TV-BKL and SR-GDN from ten sets of clinical scans acquired on a 3T scanner. The results show that both methods led to small standard deviations of estimations. Figure 6(c) shows the spectrum magnitudes of the handcraft Gaussian filter and the mean blur kernels estimated by TV-BKL and SR-GND. In comparison to the simulation results shown in Fig. 4, the results show that our approach (SR-GND) estimated the slice profile as a close approximation to a boxcar function.

We evaluated the average edge strength (AES) of the HR reconstructions obtained from our approach and the three baselines on the ten sets of clinical scans. The results, as shown in Table 1, suggests that our approach, SR-GDN achieved the highest AES. TV-BKL yielded the lowest AES as the noise in the reconstructions led to blurry edges.

Figure 7 shows the qualitative assessment of our approach and the three baselines in the slices of representative clinical data. The results show that SR-GDN offered high-quality brain MRI with the clinical LR data. GGR smoothed the images excessively so missed some anatomical details, while TV and TV-BKL generated noisy images and caused artifacts as shown in the coronal image.

### Experiment 3: Impact on Deconvolution-Based Algorithms.

We replaced the handcrafted Gaussian filter used in TV and GRR with our estimated blur kernels and ran the two methods on the HCP dataset. The increased PSNR/SSIM were respectively: TV = 13.90%/3.15% and GGR = 3.35%/2.44%. The results show that our estimated blur kernels led to improved deconvolution-based algorithms.

## 4 Discussion

We have developed a deconvolutional technique that enabled accurately estimating blur kernels from the image data itself. We have designed a deep architecture that utilizes an adversarial scheme with a generative neural network against its degradation counterparts. This design has been demonstrated to allow for the SRR tailored to an individual subject. We have thoroughly assessed the accuracy of our approach on a simulated dataset. We have successfully applied our approach to ten pediatric patients, and have achieved high-quality brain MRI at an isotropic resolution of 0.125 cubic mm with six minutes of imaging time. Experimental results have shown that our approach achieved superior SRR results as compared to state-of-the-art deconvolution-based methods, while in parallel, at substantially reduced imaging time in comparison to direct HR acquisitions.

## Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

## Acknowledgments

This work was supported in part by the National Institutes of Health (NIH) under grants R01 NS079788, R01 EB019483, R01 EB018988, R01 NS106030, R01 EB031849, IDDCR U54 HD090255, S10OD025111; a research grant from the Boston Children's Hospital Translational Research Program; a Technological Innovations in Neuroscience Award from the McKnight Foundation; a research grant from the Thrasher Research Fund; and a pilot grant from National Multiple Sclerosis Society under Award Number PP-1905-34002.

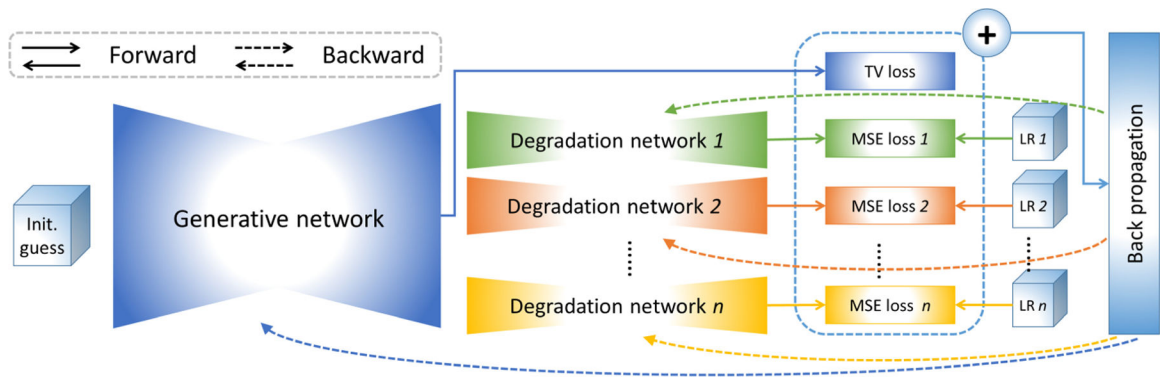
## References

1. Afacan O, et al. : Evaluation of motion and its effect on brain magnetic resonance image quality in children. *Pediatr. Radiol* 46(12), 1728–1735 (2016). [PubMed: 27488508]
2. Brown RW, Cheng YCN, Haacke EM, Thompson MR, Venkatesan R: *Magnetic Resonance Imaging: Physical Principles and Sequence Design*, 2nd edn. Wiley, New York (2014)
3. Chan T, Wong CK: Total variation blind deconvolution. *IEEE Trans. Image Process* 7(3), 370–5 (1998) [PubMed: 18276257]
4. Cherukuri V, Guo T, Schiff SJ, Monga V: Deep MR brain image super-resolution using spatio-structural priors. *IEEE Trans. Image Process* 29, 1368–1383 (2020)
5. Essen D, Smith S, Barch D, Behrens TEJ, Yacoub E, Ugurbil K: The Wu-Minn human connectome project: an overview. *NeuroImage* 80, 62–79 (2013) [PubMed: 23684880]
6. Gao H, Tao X, Shen X, Jia J: Dynamic scene deblurring with parameter selective sharing and nested skip connections. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 3843–3851 (2019)
7. Gholipour A, Estroff JA, Warfield SK: Robust super-resolution volume reconstruction from slice acquisitions: application to fetal brain MRI. *IEEE Trans. Med. Imag* 29(10), 1739–1758 (2010)
8. Gholipour A, et al. : Super-resolution reconstruction in frequency, image, and wavelet domains to reduce through-plane partial voluming in MRI. *Med. Phys* 42(12), 6919–6932 (2015) [PubMed: 26632048]
9. Greenspan H: Super-resolution in medical imaging. *Comput. J* 52(1), 43–63 (2009)
10. Greenspan H, Oz G, Kiryati N, Peled S: MRI inter-slice reconstruction using super-resolution. *Magn. Reson. Imag* 20(5), 437–446 (2002)
11. Gudbjartsson H, Patz S: The Rician distribution of noisy MRI data. *Magn. Reson. Med* 34, 910–914 (1995) [PubMed: 8598820]
12. Kingma DP, Ba J: Adam: a method for stochastic optimization. In: *International Conference on Learning Representations (ICLR)* (2015)
13. Paszke A, et al. : Automatic differentiation in pytorch. In: *Neural Information Processing Systems (NIPS) Workshop* (2017)
14. Peled S, Yeshurun Y: Super-resolution in MRI - Perhaps sometimes. *Magn. Reson. Med* 48, 409 (2002)
15. Pham C, Ducournau A, Fablet R, Rousseau F: Brain MRI super-resolution using deep 3D convolutional networks. In: *International Symposium on Biomedical Imaging*, pp. 197–200 (2017)
16. Plenge E, et al. : Super-resolution methods in MRI: can they improve the trade-off between resolution, signal-to-noise ratio, and acquisition time? *Magn. Reson. Med* 68, 1983–1993 (2012) [PubMed: 22298247]
17. Rudin L, Osher S, Fatemi E: Nonlinear total variation based noise removal algorithms. *Physica D Nonlinear Phenom* 60, 259–268 (1992)
18. Scheffler K: Super resolution in MRI? *Magn. Reson. Med* 48, 408 (2002) [PubMed: 12210953]
19. Scherrer B, Gholipour A, Warfield SK: Super-resolution reconstruction to increase the spatial resolution of diffusion weighted images from orthogonal anisotropic acquisitions. *Med. Image Anal* 16(7), 1465–1476 (2012) [PubMed: 22770597]
20. Sui Y, Afacan O, Gholipour A, Warfield SK: Isotropic MRI super-resolution reconstruction with multi-scale gradient field prior. In: Shen D, et al. (eds.) *MIC-CAI 2019. LNCS*, vol. 11766, pp. 3–11. Springer, Cham (2019).



21. Sui Y, Afacan O, Gholipour A, Warfield SK: Learning a gradient guidance for spatially isotropic MRI super-resolution reconstruction. In: Martel AL, et al. (eds.) MICCAI 2020. LNCS, vol. 12262, pp. 136–146. Springer, Cham (2020).
22. Sui Y, Afacan O, Gholipour A, Warfield SK: SLIMM: slice localization integrated MRI monitoring. *NeuroImage* 223(117280), 1–16 (2020)
23. Sui Y, Afacan O, Gholipour A, Warfield SK: Fast and high-resolution neonatal brain MRI through super-resolution reconstruction from acquisitions with var-able slice selection direction. *Front. Neurosci* 15(636268), 1–15 (2021)
24. Tourbier S, Bresson X, Hagmann P, Thiran J, Meuli R, Cuadra M: An efficient total variation algorithm for super-resolution in fetal brain MRI with adaptive regularization. *Neuroimage* 118, 584–597 (2015)
25. Ulyanov D, Vedaldi A, Lempitsky V: Deep image prior. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 9446–9454 (2018)
26. Van Reeth E, Tham IW, Tan CH, Poh CL: Super-resolution in magnetic resonance imaging: a review. *Concepts Magn. Reson* 40(6), 306–325 (2012)
27. Wang Z, Bovik AC, Sheikh HR, Simoncelli EP: Image quality assessment: from error measurement to structural similarity. *IEEE Trans. Image Process* 13(1), 600–612 (2004) [PubMed: 15376593]
28. Xue X, Wang Y, Li J, Jiao Z, Ren Z, Gao X: Progressive sub-band residual-learning network for MR image super-resolution. *IEEE J. Biomed. Health Inform* 24(2), 377–386 (2020) [PubMed: 31603805]
29. Zaca D, Hasson U, Minati L, Jovicich J: A method for retrospective estimation of natural head movement during structural MRI. *J. Magn. Reson. Imag* 48(4), 927–937 (2018)
30. Zhao X, Zhang Y, Zhang T, Zou X: Channel splitting network for single MR image super-resolution. *IEEE Trans. Image Process* 28(11), 5649–5662 (2019) [PubMed: 31217110]
31. Çiçek Ö, Abdulkadir A, Lienkamp SS, Brox T, Ronneberger O: 3D U-Net: learning dense volumetric segmentation from sparse annotation. In: Ourselin S, Joskowicz L, Sabuncu MR, Unal G, Wells W (eds.) MICCAI 2016. LNCS, vol. 9901, pp. 424–432. Springer, Cham (2016).





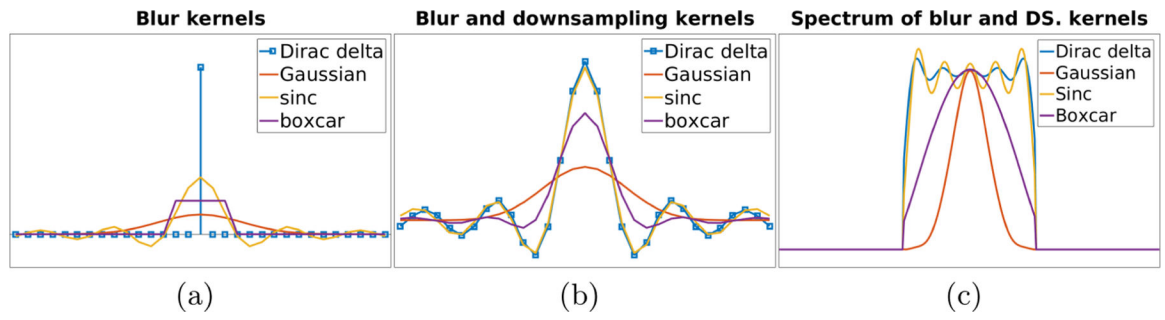
**Fig. 1.** Architecture of our proposed approach to generative degradation learning.

Author Manuscript


Author Manuscript

Author Manuscript

Author Manuscript

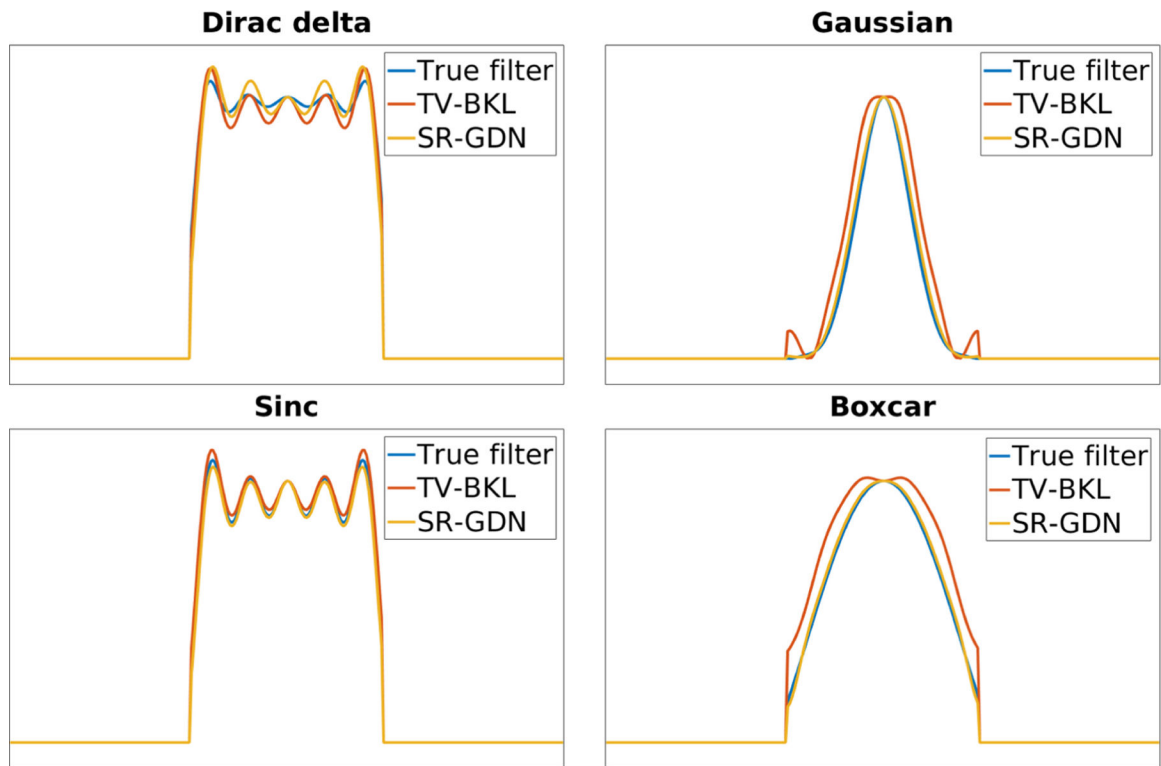


**Fig. 2.** Illustrations of the mock slice profiles (a), generated blur kernels (b), and spectrum magnitudes of these kernels (c).

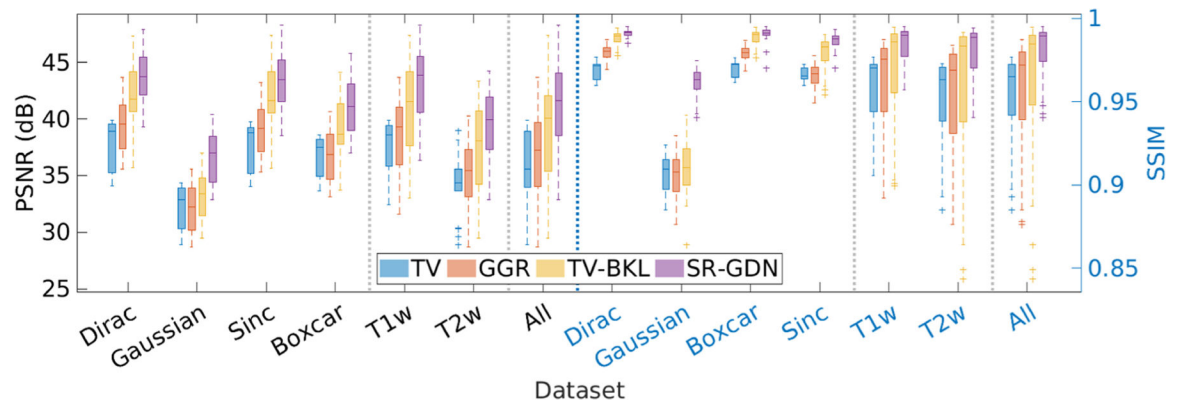


**Fig. 3.**

Estimation errors in the blur kernels obtained from the two methods. Our methods, SR-GDN, considerably outperformed TV-BKL on the HCP dataset.



**Fig. 4.** Spectrum of true and estimated blur kernels from SR-GDN and TV-BKL on four representation simulations. Our estimates (SR-GDN) closely followed the true kernels and offered higher accuracy than TV-BKL with all types of kernels.




**Fig. 5.** Quantitative assessment of the four SRR methods in terms of PSNR and SSIM.

Author Manuscript


Author Manuscript

Author Manuscript

Author Manuscript



**Fig. 6.** Mean and standard deviation of the blur kernels estimated by (a) TV-BKL and (b) SR-GDN from ten sets of clinical scans. (c) Spectrum magnitudes of the handcraft Gaussian filter and the mean blur kernels estimated by TV-BKL and SR-GDN.



**Fig. 7.**  
Qualitative assessment of our approach (SR-GDN) and the three baselines in the slices of representative clinical data.



**Table 1.**

Average edge strength obtained from our approach and the three baselines on the ten sets of clinical scans.

<b>TV</b>	<b>GGR</b>	<b>TV-BKL</b>	<b>SR-GDN</b>
$2.73 \pm 0.72$	$3.12 \pm 0.90$	$1.78 \pm 0.50$	$3.49 \pm 0.97$

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript