LEARNING FROM A HANDFUL VOLUMES: MRI RESOLUTION ENHANCEMENT WITH VOLUMETRIC SUPER-RESOLUTION FORESTS

Aline Sindel[1] Katharina Breininger[1] Johannes Käßer[2] Andreas Hess[2] Andreas Maier[1,3] Thomas Köhler[1,3,4]

[1]Pattern Recognition Lab, FAU Erlangen-Nürnberg, Germany
[2]Institute for Experimental & Clinical Pharmacology & Toxicology, FAU Erlangen-Nürnberg, Germany
[3]Erlangen Graduate School in Advanced Optical Technologies (SAOT), FAU Erlangen-Nürnberg, Germany
[4]e.solutions GmbH, Erlangen, Germany
`https://github.com/asindel/VSRF`

ABSTRACT

Magnetic resonance imaging (MRI) enables 3-D imaging of anatomical structures. However, the acquisition of MR volumes with high spatial resolution leads to long scan times. To this end, we propose volumetric super-resolution forests (VSRF) to enhance MRI resolution retrospectively. Our method learns a locally linear mapping between low-resolution and high-resolution volumetric image patches by employing random forest regression. We customize features suitable for volumetric MRI to train the random forest and propose a median tree ensemble for robust regression. VSRF outperforms state-of-the-art example-based super-resolution in terms of image quality and efficiency for model training and inference on different MRI datasets. It is also superior to unsupervised methods with just a handful or even a single volume to assemble training data.

Index Terms— Super-resolution, random forests, MRI

1. INTRODUCTION

High-resolution magnetic resonance imaging (MRI) allows the visualization of delicate anatomical structures in-vivo, which is crucial to support early detection of pathologies and to enable an accurate prediction of their size and composition. However, high-resolution MRI for this task requires long acquisition times, which can lead to stress and discomfort for the imaged subject. To reduce acquisition times while retaining high spatial resolution, resolution enhancement can be applied during data acquisition, e.g. using zero-filling, or retrospectively by means of super-resolution (SR) [1].

SR estimates high-resolution (HR) images from single or sets of low-resolution (LR) images. *Multi image* SR [2, 3, 4] can effectively enhance the spatial resolution but multiple acquisitions increase scan times. *Single image* (SISR) methods are promising alternatives. *Reconstruction-based* SISR methods [5, 6] are based on a regularized optimization problem to enforce the downsampled version of the predicted HR result to be close to the LR image. *Example-based* methods estimate a HR image from a single LR image based on pairs of LR/HR examples of an external database. In this area, dictionary-based approaches [7, 8, 9, 10] build on sparse representations of image patches. Regression-based approaches are more effective as they avoid time-consuming sparse coding. These include tree-based regression [11] on a patch level as well as deep learning [12, 13, 14] to infer end-to-end mappings. Another class of SISR, *self* SR [15, 16], estimates the HR image by internally learning from patches of different scales without external databases.
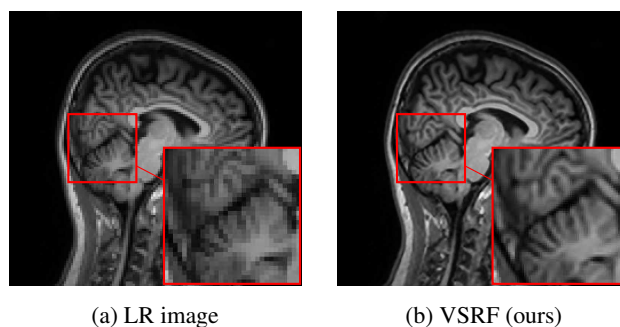

(a) LR image   (b) VSRF (ours)

Fig. 1: Our *volumetric super-resolution forests* (VSRF) facilitate high-resolution MRI while retaining a low scan time as illustrated by one sagittal slice of the Kirby 21 MRI dataset [17] (SR factor 2).

SISR techniques in 3-D MRI can be either used on *slice* or *volume* level. Slice-wise methods enhance the resolution within one plane but ignore the coherence between adjacent slices of volumetric data. Volumetric methods enable simultaneous resolution enhancement in all directions. For instance, Manjón et al. [18] applied an iterative patch-based non-local reconstruction scheme based on self-similarity and the 3-D non-local means filter. Jog et al. [19] build up on anchored neighborhood regression [9] to perform self SR. A closely related field to MRI SR, image synthesis, uses dictionaries [20] or random forests [21] to super-resolve MR volumes by utilizing additional information of a HR volume with another MR contrast of the same subject. Alexander et al. [22] implemented a random forest for SR in diffusion MRI with cubic patches and features customized for diffusion tensor images (DTIs). Yoldemir et al. [23] applied dictionary learning to diffusion-weighted 3-D images for volumetric SR. Recently, Tanno et al. [24] integrated an uncertainty modeling with a 3-D convolutional neural network (CNN) on DTIs. Bahrami et al. [25] used a CNN with anatomical and appearance features to non-linearly map 3T to 7T MR volumes. In a further work [26], they estimated a 7T-like volume with a random forest based on 3T and 7T patches and enhanced the 7T-like volume with dictionary learning. For brain MRI, Rueda et al. [27] applied dictionary learning for 3-D SR. More recently, Pham et al. [28] and Chen et al. [29] proposed deep learning by extending 2-D CNNs for natural images to 3-D CNNs. CNNs avoid the need for hand-crafted features and enable end-to-end learning. This leads to state-of-the art results but requires large amounts of training data that need to match to the desired application to be effective [30]. In contrast, random forests yield robust results with only a small amount of training data and are

very fast to train even without sophisticated hardware.

In this paper, we propose *volumetric SR forests* (VSRF) for example-based resolution enhancement in MRI. Our method employs random forest regression and generalizes SR forests (SRF) [11] that have been originally introduced for natural images to volumetric data. We train the random forest with overlapping 3-D patches and new *customized features*. For robust regression, we propose a *median ensemble model* to obtain the forest prediction. Our proposed VSRF can effectively improve the spatial resolution in MRI, see Fig. 1. It facilitates computationally efficient training and inference even on a CPU. Moreover, it is effective for example-based SR even with a limited amount of training data making it an attractive tool for practical applications.

2. VOLUMETRIC SUPER-RESOLUTION FORESTS

Our proposed VSRF builds on random forest regression [11] to learn a locally linear mapping between LR and HR patches from example data. We introduce training, inference and features for this model.

2.1. Random Forest Training

A random forest consists of a set of decision trees which are independently trained on a training set of N patch pairs $\{\mathbf{x}_L^n, \mathbf{x}_H^n\}_{n=1}^N$, where $\mathbf{x}_L \in \mathbb{R}^{D_L}$ and $\mathbf{x}_H \in \mathbb{R}^{D_H}$ are feature vectors representing the LR and HR patches. The feature vectors $\mathbf{x}_L^n$ and $\mathbf{x}_H^n$ are stacked into matrices $\mathbf{X}_L \in \mathbb{R}^{D_L \times N}$ and $\mathbf{X}_H \in \mathbb{R}^{D_H \times N}$ [11]. The tree structure is learned by recursively dividing the data space into disjoint subsets until the maximum tree depth or minimum number of feature vector pairs in a node is reached and a leaf node is created. The splits are performed based on a binary decision made at each internal node using a pair-wise difference splitting function [11, 31]:

$$h(\mathbf{x}_L, \boldsymbol{\theta}) = \begin{cases} 1 & \mathbf{x}_L[\varphi_1] - \mathbf{x}_L[\varphi_2] < \tau, \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

which compares the difference of two randomly selected feature dimensions $\varphi_1, \varphi_2 \in \{1, \ldots, D_L\}$ from the LR feature vector $\mathbf{x}_L$ to a threshold τ. The parameters $\boldsymbol{\theta} = \{\varphi_1, \varphi_2, \tau\}$ for the splitting functions are estimated by evaluating a quality measure. We use node optimization for finding the splitting parameters, i.e. we do not evaluate the quality measure on all training samples of the node but we randomly subsample the data of that node [11]. Thus, we reduce computation time and increase the variation between the trees. A common choice for the quality measure for regression forests is reduction-in-variance based on information gain [26, 31]. We use a modified quality measure defined by Schulter et al. [11] that operates in both the low and high-resolution space according to:

$$Q(\mathbf{X}_H, \mathbf{X}_L, \boldsymbol{\theta}) = \sum_{i \in \{left, right\}} N_i \cdot E(\mathbf{X}_H^i, \mathbf{X}_L^i), \quad (2)$$

where $\mathbf{X}_L^i$ and $\mathbf{X}_H^i$ contain the N_i LR and HR feature vectors of the left and right subsets according to $\boldsymbol{\theta}$. The optimization of the quality measure selects the splits for which the variance of HR feature vectors and variance of LR feature vectors in the subsets is minimal. The variance for both domains [11] is given by:

$$E(\mathbf{X}_H^i, \mathbf{X}_L^i) = \frac{1}{N_i} \sum_{n=1}^{N_i} (\|\mathbf{x}_H^n - \bar{\mathbf{x}}_H^n\|_2^2 + \kappa \cdot \|\mathbf{x}_L^n - \bar{\mathbf{x}}_L^n\|_2^2), \quad (3)$$

where $\bar{\mathbf{x}}$ is the mean over all samples and κ is a hyper-parameter to control the influence of the LR variance.

In the leaf nodes, we learn locally linear mappings from LR feature vectors to HR feature vectors using the feature vector pairs that reach these specific nodes. We determine a mapping $\hat{\mathbf{W}}_l$ for the leaf l according to the least squares problem [11]:

$$\hat{\mathbf{W}}_l = \underset{\mathbf{W}_l}{\arg\min} \sum_{n=1}^{N_l} \|\mathbf{x}_H^n - \mathbf{W}_l \cdot \mathbf{x}_L^n\|_2^2. \quad (4)$$

To yield a more stable solution, we use ridge regression:

$$\hat{\mathbf{W}}_l = \underset{\mathbf{W}_l}{\arg\min} \|\mathbf{X}_H^l - \mathbf{W}_l \mathbf{X}_L^l\|_2^2 + \lambda \|\mathbf{W}_l\|_2^2, \quad (5)$$

where λ is the regularization parameter. The estimate for the linear mapping can be solved in closed-form [11]:

$$\hat{\mathbf{W}}_l^\top = (\mathbf{X}_L^{l\top} \mathbf{X}_L^l + \lambda \mathbf{I})^{-1} \mathbf{X}_L^{l\top} \cdot \mathbf{X}_H^l. \quad (6)$$

2.2. Random Forest Inference

In random forest inference the LR feature vectors of the test set are sent through all trees resulting in one prediction of the mapping function for each LR feature vector and each tree. The estimated predictions are then combined to a single HR patch by the forest ensemble model. A common approach for ensembles is to average the prediction of all trees. For SR this means computing an element-wise average of the predicted HR patches. To account for outliers of the predicted values of the trees, we employ a median ensemble model which is more robust with regards to these issues and computes the median for each component of the predicted HR patches. The predicted HR information is then added to the tricubically upsampled LR volume. Since we extracted overlapping patches, we reconstruct the final volume by averaging the overlapping voxels.

2.3. Features and Patch Extraction

In the proposed VSRF method, LR volumes are processed in two steps: 1) The LR volume $\mathbf{V}_L$ is upscaled to the size of the target HR volume by tricubic interpolation resulting in the volume $\tilde{\mathbf{V}}_L$. 2) The missing high-frequency information is predicted by the random forest and added to $\tilde{\mathbf{V}}_L$. In order to perform 3-D SR, we extract overlapping $n_x \times n_y \times n_z$ patches from $\tilde{\mathbf{V}}_L$. We extract corresponding HR patches of equal size from the difference volume $\mathbf{V}_H - \tilde{\mathbf{V}}_L$ that contains the missing high-frequency information of $\tilde{\mathbf{V}}_L$ compared to the HR volume $\mathbf{V}_H$.

Based on the LR patches, we compute a set of ten different features to which we refer as DevEdge in the following: Partial first- and second-order derivatives in all three directions (Dev), edge orientation in all three directions and edge magnitude. First- and second-order derivatives are used in the SR framework of [9] for 2-D images, which we extend to 3-D by adding a third component in z-direction. Edge magnitude M and edge orientation ϕ are computed based on the first-order derivatives D_i with $i \in \{x, y, z\}$:

$$M = \sqrt{D_x^2 + D_y^2 + D_z^2}, \quad \phi_{xy} = \arctan\left(\frac{D_y}{D_x}\right), \quad (7)$$

the edge orientations ϕ_{zx} and ϕ_{zy} are computed accordingly. Since edge orientation is very sensitive to noise, we convolve $\tilde{\mathbf{V}}_L$ with a Gaussian filter with standard deviation σ before we compute the derivatives for edge orientation and magnitude. To accelerate the training process, a PCA dimensionality reduction is applied to the feature vectors like in [9]. These feature vectors together with the vectorized HR patches are then used as input for the random forest.

Table 1: Mean PSNR (dB) and SSIM on the mouse brain dataset (5 volumes) and on the Kirby 21 human brain dataset [17] (30 volumes). We compared the proposed volumetric SR forest (VSRF) to different state-of-the-art SR methods (SR factor 2).

Dataset	Tricubic		2-D SRF [11]		Psd. 3-D SRCNN		Psd. 3-D SRF		NLMU [18]		VANR		VA+		VSRF (ours)	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
Mouse Brain	34.94	0.9637	37.31	0.9752	36.82	0.9680	38.63	0.9781	36.94	0.9721	37.69	0.9750	38.75	0.9779	**39.46**	**0.9804**
Kirby Brain	34.84	0.9502	35.58	0.9611	36.10	0.9643	36.48	0.9659	36.58	0.9662	35.59	0.9605	36.06	0.9650	**37.15**	**0.9701**

3. EXPERIMENTS AND RESULTS

We evaluated the performance of VSRF with a mouse and a human brain MRI dataset. The mouse brain MRI dataset consisted of 21 isotropic high-resolution volumes[1]. The MR volumes were acquired with a 4.7T MR scanner (Bruker BioSpec 47/40) with $0.05 \times 0.05 \times 0.05 \, \text{mm}^3$ resolution, which were measured in-situ after transcardial perfusion with ProHance$^{\text{TM}}$ with a T1-weighted 3-D FLASH sequence ($TE = 4.6 \, \text{ms}$, $TR = 16.3 \, \text{ms}$, field of view (FOV): $17 \, \text{mm} \times 15 \, \text{mm} \times 8 \, \text{mm}$, matrix: $340 \times 300 \times 160$ voxels, flip angle $= 20°$, $NEX = 14$, total scan time $= 3 \, \text{h} \, 2 \, \text{min} \, 37\text{s}$). We manually segmented the mouse brain and cropped the volumes to $320 \times 210 \times 140$ voxels. The LR volumes were generated by tricubic downscaling of the HR volumes by an isotropic scaling factor of 2. For LR volumes the acquisition time would be reduced substantially by a factor of 5.58 to $32 \, \text{min} \, 42 \, \text{s}$ at $NEX = 10$ keeping comparable signal-to-noise ratio (SNR). We used 13 pairs of LR and HR volumes for training, three to validate our choice of parameters and five for testing. Second, we applied the SR methods to the publicly available Kirby 21 dataset [17] with human brain MR data. We used the 42 T1-weighted MPRAGE volumes with $1.0 \times 1.0 \times 1.2 \, \text{mm}^3$ resolution, which were acquired with a 3T MR scanner (Achieva, Philips Healthcare) in the sagittal plane and we cropped them to $240 \times 204 \times 256$ voxels. The LR volumes were generated by tricubic downscaling by a factor of 2. The 42 pairs of HR and LR volumes were divided into a training set of ten (No. 33-42), a validation set of two (No. 31-32) and a test set of 30 volume pairs (No. 1-30).

Comparison to State-of-the-Arts. For both datasets, we compared the performance of VSRF with tricubic upsampling, 3-D SR methods, pseudo 3-D SR methods and the 2-D SRF [11]. SRF was applied slice-wise to the volumetric MR data using 3×3 patches, $T = 30$ trees, $\kappa = 1$ and λ was estimated automatically (according to the source code of [11]) based on the condition number of the least squares problem. The parameter settings for our VSRF are: $T = 30$ trees, $3 \times 3 \times 3$ patches, $\kappa = 1$ and automatic λ and $\sigma = 1$ for the edge orientation features. For 3-D SR, we extended the 2-D anchored neighborhood regression (ANR) [9] and adjusted anchored neighborhood regression (A+) [10] by employing 3-D patches and utilizing the same features as for VSRF. We refer to these volumetric methods as VANR and VA+. For VANR and VA+ we used $3 \times 3 \times 3$ patches, a dictionary size of 2048 and neighborhoods of 16 atoms for VANR and 2048 samples for VA+. We directly utilized all training samples for the computation of VA+ regressors instead of augmenting the samples by means of a scaling pyramid and then restricting the number to 5 million like in [10]. Also we applied the non-local MRI upsampling (NLMU) [18] as a 3-D SR method. In addition, we extended the 2-D SRF [11] and 2-D SRCNN [12] to pseudo 3-D SRF and pseudo 3-D SRCNN. Following prior work [28], the pseudo 3-D methods were constructed by averaging the three slice-wise SRFs

or SRCNNs in axial, sagittal and coronal view of the MR volume. We trained the 2-D SRCNN from scratch with an augmented MRI dataset over 15 million iterations (1383 epochs for the mouse and 1633 epochs for the Kirby 21 dataset) by feeding the SRCNN with slices of the axial, sagittal and coronal view of the volume.

The SR results on the mouse and Kirby 21 dataset were evaluated using the peak signal-to-noise ratio (PSNR) and the structural similarity (SSIM) [32], see Table 1. VSRF achieved the highest PSNR and SSIM of all tested methods on both datasets. On the mouse dataset VSRF achieved a gain in PSNR (SSIM) of $0.71 \, \text{dB}$ (0.0025) compared to VA+ and of $0.83 \, \text{dB}$ (0.0023) compared to pseudo 3-D SRF (next best results). On the Kirby 21 dataset VSRF outperformed NLMU (second best) with a $0.57 \, \text{dB}$ (0.0039) higher PSNR (SSIM). Compared to tricubic upsampling VSRF increased the PSNR (SSIM) by $4.52 \, \text{dB}$ (0.0167) on the mouse and by $2.32 \, \text{dB}$ (0.0199) on the Kirby dataset.

In Fig. 2 regions of interest (ROIs) from the sagittal view of a mouse brain and in Fig. 3 ROIs from the coronal view of a human brain from the Kirby 21 dataset are visualized. The SR methods are successful in learning and recovering high-frequency information lost in the tricubically upsampled volumes (Fig. 2b and Fig. 3b). The cerebellar sub-structure (arbor vitae) of the mouse brain is considerably sharper in the VSRF ROI (Fig. 2i). The other 3-D SR results (Fig. 2f,g,h) are blurrier than VSRF with less distinct contours, e.g. VA+ shows some artifacts around the arbor vitae. Some fine structures of the human cerebellum (Fig. 3g) which are barely visible using tricubic upsampling (Fig. 3b) regain sharpness to some extent with the SR methods. VSRF (Fig. 3f) generates a more precise visualization compared to the other 3-D methods (Fig. 3c,d,e) due to artifacts of the VA+ and the fuzziness of the VANR and NLMU results. VSRF seems to better utilize the features to learn the characteristics of the MRI dataset than VANR and VA+. Further, the choice of the median ensemble model added additional stability against outliers. The pseudo 3-D methods (Fig. 2d,e) are also blurrier than VSRF, especially the SRCNN that over-emphasizes the bright area around the arbor vitae. We observe an increase in image quality from the 2-D SRF (Fig. 2c) to the pseudo 3-D SRF to VSRF. We partly attribute the improvement of VSRF in all three directions compared to SRF and the pseudo 3-D methods to the use of a cubic neighborhood as it better exploits the spacial information contained in the volumes, while the 2-D SRF only captures in-plane relationships. The pseudo 3-D methods average the results of the 2-D methods of the three views, hence recover more information than the 2-D methods but are blurrier than the 3-D methods.

Influence of Parameters for VSRF. We evaluated the influence on the image quality by the number of training volumes for VSRF, see Fig. 4. We observed a rapid increase in performance for up to five volumes. More volumes led to minor improvements reaching the highest PSNR and SSIM values for 13 volumes. Even by using only one training volume, VSRF outperformed tricubic upsampling and NLMU, which do not require training data. Hence, VSRF indicated to be relatively robust with a small amount of training data.

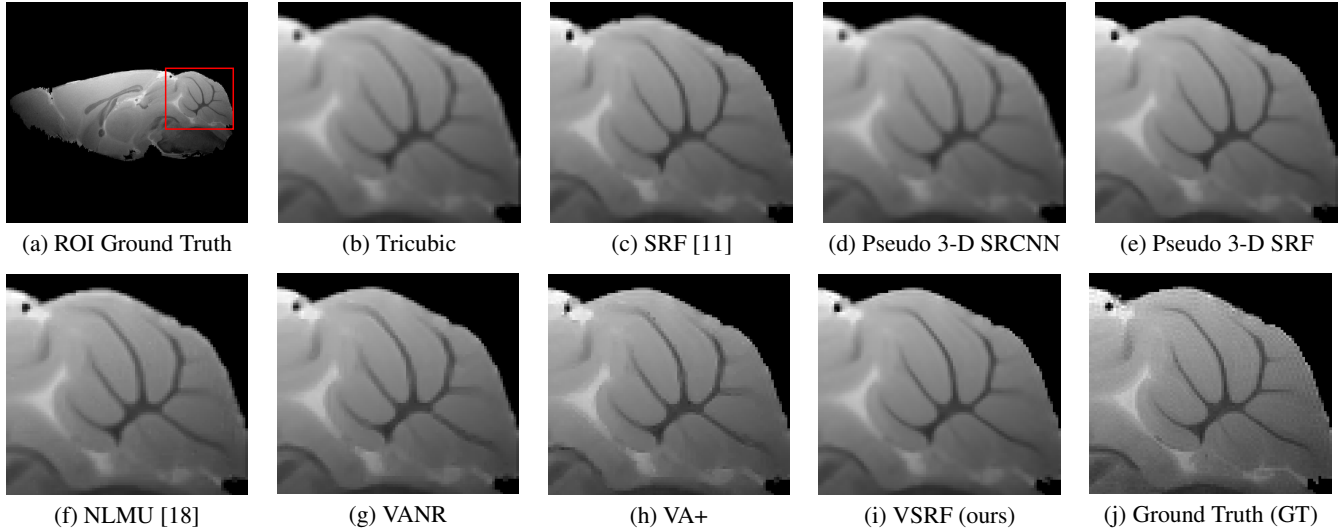[1]All animal experiments were approved by the local ethic committee (Regierung von Unterfranken, Würzburg, Germany)

(a) ROI Ground Truth   (b) Tricubic   (c) SRF [11]   (d) Pseudo 3-D SRCNN   (e) Pseudo 3-D SRF

(f) NLMU [18]   (g) VANR   (h) VA+   (i) VSRF (ours)   (j) Ground Truth (GT)

Fig. 2: Visual comparison of SR results for one sagittal slice of the mouse brain MRI dataset (SR factor 2).

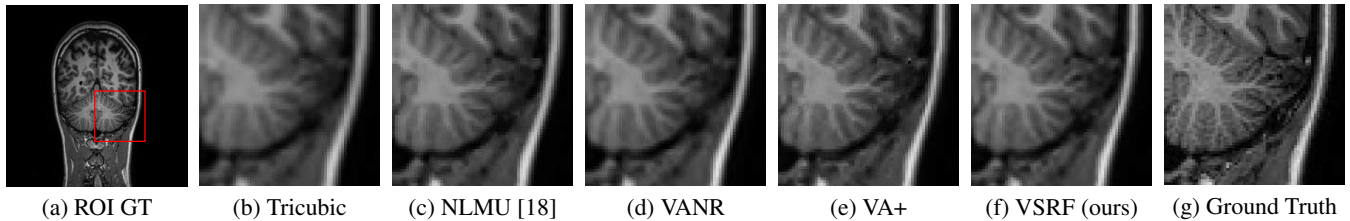(a) ROI GT   (b) Tricubic   (c) NLMU [18]   (d) VANR   (e) VA+   (f) VSRF (ours)   (g) Ground Truth

Fig. 3: Visual comparison of SR results for one coronal slice of the Kirby 21 MRI dataset [17] (SR factor 2).
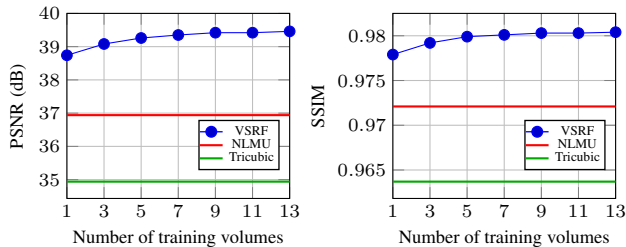

Fig. 4: Influence of the number of training volumes for VSRF on the mouse brain dataset (SR factor 2) evaluated with mean PSNR (left) and SSIM (right). NLMU [18] and tricubic upsampling results are plotted for comparison.

Fig. 5 showed the clear advantage of the median ensemble for VSRF compared to the average ensemble regarding PSNR and SSIM, partly due to patch artifacts caused by the average ensemble. The feature set DevEdge yielded additional improvement compared to Dev for the median ensemble.

Computation Time. The VSRF is very fast in training and inference. The inference for the mouse brain MRI dataset required less than 1 min per volume and the training of the 13 volumes took about 1 h (or about 20 min in parallel) with an Intel i7 CPU 3.4 GHz. VANR and VA+ with about 20 s per volume were even faster in inference but noticeable slower in training due to the dictionary learning (VANR required 5 h 15 min and VA+ 7 h 32 min). The execution time for NLMU took about 90 s per volume. SRCNN required around 3 days for training on a GPU GeForce GTX 1080 and
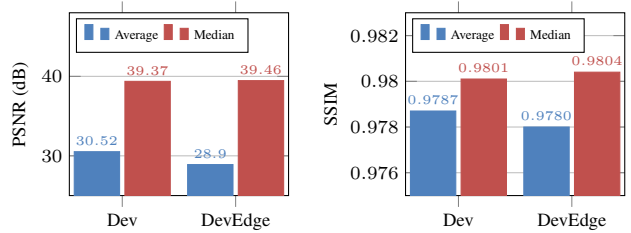

Fig. 5: Influence of the ensemble model (average and median) and features (Dev and DevEdge) for VSRF on the mouse brain dataset (SR factor 2) evaluated with mean PSNR (left) and SSIM (right).

pseudo 3-D SRCNN about 10 min per volume for inference using the MATLAB reference implementation of [12].

4. CONCLUSION

We presented a volumetric SR method for brain MRI based on random forests that learn mappings of 3-D LR to HR patches. In the experiments with the mouse brain and the Kirby 21 human brain MRI datasets, our VSRF approach demonstrated visually and quantitatively an improvement in image quality compared to the state-of-the-art and achieved fast training and inference performance with a small amount of training data. The proposed approach for MRI resolution enhancement may be utilized to remarkably reduce MRI acquisition time with smaller loss of image quality, which makes adaption into clinical workflows appealing.

5. REFERENCES

[1] E. Van Reeth, I. W. K. Tham, C. H. Tan, and C. L. Poh, "Super-resolution in magnetic resonance imaging: A review," *Concepts in Magnetic Resonance Part A*, vol. 40A, no. 6, pp. 306–325, 2012.

[2] T. Köhler, X. Huang, F. Schebesch, A. Aichert, A. Maier, and J. Hornegger, "Robust Multiframe Super-Resolution Employing Iteratively Re-Weighted Minimization," *IEEE Trans Comput Imaging*, vol. 2, no. 1, pp. 42–58, 2016.

[3] C. Bercea, A. Maier, and T. Köhler, "Confidence-aware Levenberg-Marquardt optimization for joint motion estimation and super-resolution," in *Proc. IEEE ICIP 2016*, 2016, pp. 1136–1140.

[4] T. Köhler, M. Bätz, F. Naderi, A. Kaup, A. K. Maier, and C. Riess, "Benchmarking super-resolution algorithms on real data," *arXiv preprint arXiv:1709.04881*, 2017.

[5] S. Dai, M. Han, Y. Wu, and Y. Gong, "Bilateral back-projection for single image super resolution," in *Proc. IEEE ICME 2007*, 2007, pp. 1039–1042.

[6] M. N. Bareja and C. K. Modi, "An effective iterative back projection based single image super resolution approach," in *Proc. IEEE CSNT 2012*, 2012, pp. 95–99.

[7] J. Yang, J. Wright, T. S. Huang, and Y. Ma, "Image super-resolution via sparse representation," *IEEE Trans. Img. Proc.*, vol. 19, no. 11, pp. 2861–2873, 2010.

[8] R. Zeyde, M. Elad, and M. Protter, "On single image scale-up using sparse-representations," in *Curves and Surfaces 2010*, 2012, pp. 711–730.

[9] R. Timofte, V. De Smet, and L. Van Gool, "Anchored neighborhood regression for fast example-based super-resolution," in *Proc. IEEE ICCV 2013*, 2013, pp. 1920–1927.

[10] R. Timofte, V. De Smet, and L. Van Gool, "A+: Adjusted anchored neighborhood regression for fast super-resolution," in *Proc. ACCV 2014*, 2015, pp. 111–126.

[11] S. Schulter, C. Leistner, and H. Bischof, "Fast and accurate image upscaling with super-resolution forests," in *Proc. IEEE CVPR 2015*, 2015, pp. 3791–3799.

[12] C. Dong, C. C. Loy, K. He, and X. Tang, "Learning a deep convolutional network for image super-resolution," in *Proc. ECCV 2014*, 2014, pp. 184–199.

[13] J. Kim, J. K. Lee, and K. M. Lee, "Accurate Image Super-Resolution Using Very Deep Convolutional Networks," in *Proc. IEEE CVPR 2016*, 2016, pp. 1646–1654.

[14] B. Lim, S. Son, H. Kim, S. Nah, and K. M. Lee, "Enhanced deep residual networks for single image super-resolution," in *Proc. IEEE CVPRW 2017*, 2017, pp. 1132–1140.

[15] D. Glasner, S. Bagon, and M. Irani, "Super-resolution from a single image," in *Proc. IEEE ICCV 2009*, 2009, pp. 349–356.

[16] G. Freedman and R. Fattal, "Image and video upscaling from local self-examples," *ACM Trans. Graph.*, vol. 28, no. 3, pp. 1–10, 2010.

[17] B. A. Landman, A. J. Huang, A. Gifford, D. S. Vikram, I. A. L. Lim, J. A. D. Farrell, J. A. Bogovic, J. Hua, M. Chen, S. Jarso, S. A. Smith, S. Joel, S. Mori, J. J. Pekar, P. B. Barker, J. L. Prince, and P. C. M. van Zijl, "Multi-parametric neuroimaging reproducibility: A 3-T resource study," *NeuroImage*, vol. 54, no. 4, pp. 2854 – 2866, 2011.

[18] J. V. Manjón, P. Coupé, A. Buades, V. Fonov, D. Louis Collins, and M. Robles, "Non-local MRI upsampling.," *Med. Image Anal.*, vol. 14, no. 6, pp. 784–92, 2010.

[19] A. Jog, A. Carass, and J. L. Prince, "Self super-resolution for magnetic resonance images," in *Proc. MICCAI 2016*, 2016, pp. 553–560.

[20] S. Roy, A. Carass, and J. L. Prince, "Magnetic resonance image example-based contrast synthesis," *IEEE Trans. Med. Imag.*, vol. 32, no. 12, pp. 2348–2363, 2013.

[21] A. Jog, A. Carass, and J. L. Prince, "Improving magnetic resonance resolution with supervised learning," in *Proc. IEEE ISBI 2014*, 2014, pp. 987–990.

[22] D. Alexander, D. Zikic, J. Zhang, H. Zhang, and A. Criminisi, "Image quality transfer via random forest regression: Applications in diffusion MRI," *Proc. MICCAI 2014*, vol. 17, pp. 225–32, 2014.

[23] B. Yoldemir, M. Bajammal, and R. Abugharbieh, "Dictionary based super-resolution for diffusion MRI," in *Proc. MICCAI CDMRI Workshop 2014*, 2014, pp. 203–213.

[24] R. Tanno, D. E. Worrall, A. Ghosh, E. Kaden, S. N. Sotiropoulos, A. Criminisi, and D. C. Alexander, "Bayesian image quality transfer with CNNs: Exploring uncertainty in dMRI super-resolution," in *Proc. MICCAI 2017*, 2017, pp. 611–619.

[25] K. Bahrami, F. Shi, I. Rekik, and D. Shen, "Convolutional neural network for reconstruction of 7T-like images from 3T MRI using appearance and anatomical features," in *Proc. MICCAI LABELS/DLDLM workshop 2016*, 2016, pp. 39–47.

[26] K. Bahrami, F. Shi, I. Rekik, Y. Gao, and D. Shen, "7T-guided super-resolution of 3T MRI," *Medical Physics*, vol. 44, no. 5, pp. 1661–1677, 2017.

[27] A. Rueda, N. Malpica, and E. Romero, "Single-image super-resolution of brain MR images using overcomplete dictionaries," *Med. Image Anal.*, vol. 17, no. 1, pp. 113 – 132, 2013.

[28] C. H. Pham, A. Ducournau, R. Fablet, and F. Rousseau, "Brain MRI super-resolution using deep 3D convolutional networks," in *Proc. IEEE ISBI 2017*, 2017, pp. 197–200.

[29] Y. Chen, Y. Xie, Z. Zhou, F. Shi, A. G. Christodoulou, and D. Li, "Brain MRI super resolution using 3D deep densely connected neural networks," *arXiv preprint arXiv:1801.02728*, 2018.

[30] H. S. Mousavi, T. Guo, and V. Monga, "Deep Image Super Resolution via Natural Image Priors," *arXiv preprint arXiv:1802.02721*, 2018.

[31] P. Dollár and L. Zitnick, "Structured forests for fast edge detection," in *Proc. IEEE ICCV 2013*, 2013, pp. 1841–1848.

[32] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, 2004.