MRI Super-Resolution using Multi-Channel Total Variation

Mikael Brudfors[1] (✉), Yaël Balbastre[1], Parashkev Nachev[2], and John Ashburner[1]

[1] The Wellcome Centre for Human Neuroimaging, UCL, London, UK
{mikael.brudfors.15,y.balbastre,j.ashburner}@ucl.ac.uk
[2] UCL Institute of Neurology, London, UK
p.nachev@ucl.ac.uk

Abstract. This paper presents a generative model for super-resolution in routine clinical magnetic resonance images (MRI), of arbitrary orientation and contrast. The model recasts the recovery of high resolution images as an inverse problem, in which a forward model simulates the slice-select profile of the MR scanner. The paper introduces a prior based on multi-channel total variation for MRI super-resolution. Bias-variance trade-off is handled by estimating hyper-parameters from the low resolution input scans. The model was validated on a large database of brain images. The validation showed that the model can improve brain segmentation, that it can recover anatomical information between images of different MR contrasts, and that it generalises well to the large variability present in MR images of different subjects.

Keywords: Super-resolution, Multi-channel total variation, MRI, ADMM

1 Introduction

The cost of storing data has decreased dramatically in recent decades and large databases of patient images are now contained within most hospitals. In fact, up until 2010, five billion medical imaging studies had been conducted worldwide [1]. However, although large quantities of data exist, much of the information available in that data lies latent, its potential uses yet untapped [2]. This is because hospital grade MR data commonly are of lower quality than images collected in a research context. This quality difference hampers automated neuroimaging analysis and originates from the scans being acquired in shorter times, leading to greater prevalence of different types of artefacts, in particular thick-sliced data.

The anatomical information in hospital MR images differs from images used in research studies in being distributed across a larger number of low resolution (LR) scans. A collection of MR images from a patient scanning session therefore commonly contains: (1) thick-sliced data and (2) multiple MR contrasts (e.g. T1- and T2-weighted images). For example, in a research study, one might acquire a single, time-consuming, high-resolution (HR) T1-weighted image. Whereas in a

hospital study, one might acquire several quick LR scans of varying contrasts and with thick slices. Multi-contrast MR images are acquired in order to highlight different anatomical information, in particular pathology. Hospital MR images can therefore be thought of as information distributed across multiple scans.

If the distributed information across hospital MR images could be consolidated it could be possible to achieve closer to research quality imaging using ordinary, hospital quality data. Reconstructing HR images in such a way would in turn enable large scale studies, which have been difficult to perform in a research environment because of the expense of scanning a large number of patients within a research context. A framework that could merge this distributed information would therefore be of value in, for example, research relying on learning from a population of MR images (see e.g [3–5]).

This paper proposes a generative model that given a set of patient MR images, which could be of different contrasts, reconstructs isotropic HR images, using a super-resolution technique. The paper furthermore introduces the multi-channel total variation norm as a prior for super-resolution recovery in MR images. The model additionally does not require any HR reference data, but could utilise such data if available. The incentives for choosing a generative model are: (1) the large variability in clinical MRI, which can be difficult to model in a discriminative setting; and (2), that learning methods potentially hallucinate information not actually present, which should be avoided when accurate patient diagnosis is of interest. Because of the challenge in quantifying the results of quality enhancement applied to MR images, the model was validated on thick-sliced data generated from the IXI dataset[3].

2 Background

Super-resolution (SR) methods construct HR images from several observed LR images, where each LR image transforms and samples the same HR scene. SR methods are therefore able to utilise distributed information available in collections of images, such as patient MR data. SR have been studied in MRI for numerous applications [6], and shown to improve the resolution and signal-to-noise ratio favourably compared with direct HR acquisition [7].

Most SR work in MRI have focused on images of the brain. This is because SR methods have a high dependency on precise registration of the observed LR images, and for the brain, simply a rigid alignment is sufficient. SR methods applied to brain MR scans can be classified as either single- or multi-modal. Single-modality methods combine LR images of the same MR contrasts into a single HR image (e.g. three T1-weighted scans of the same subject). Numerous methods have been proposed for single-modality SR based on e.g. the acquisition model [8, 9], image patches [10], random forests [11] and convolutional neural networks [12]. A multitude of prior terms (or regularisation methods) have also been investigated, e.g. Tikhonov [9], total variation [13] and Beltrami [14].

[3] The IXI dataset contains nearly 600 HR MR images from healthy subjects, of multiple MR contrasts. It is available from http://brain-development.org/ixi-dataset/.

Multi-modality methods have so far been less frequently explored than their single-modality counterparts. In these techniques the goal is to utilise anatomical information distributed over a number of various quality images, of different MR contrasts (e.g. two T1- and one T2-weighted scan). Manjon et al. [15] proposed a multi-modal method that used as input HR reference data and a pre-interpolated version of the LR data. HR images were then reconstructed by averaging voxels in the interpolated version of the LR data using the HR data as reference. In a paper by Rousseau [16], the main idea was to reconstruct a HR image using one LR image and an intermodality prior, which included information from another HR image. An extension to the method proposed in [15] was suggested by Jafari et al. [17], which resulted in higher interpolation accuracy. This method used a similarity metric based on the assumption that voxels that have equal distances to a strong edge are more likely to belong to the same tissue type.

3 Methods

Compared to the multi-modal SR approaches mentioned in the previous section, the model proposed here does not require any complex learning, nor any HR reference data (which significantly reduces the applicability of any SR technique in clinical practice). Instead, it relies on the definition of a generative stochastic process. This process assumes that each LR image is generated by selecting thick slices, possibly rotated and/or translated, in a HR image, with the addition of i.i.d. Gaussian noise (e.g. two LR T1- and three LR T2-weighted images are assumed observations of one HR T1- and one HR T2-weighted image, respectively). This yields a conditional probability distribution known as the data likelihood. A HR image is also assumed to result from a random process, characterised by a probability distribution known as the prior. In SR, the prior should favour images with large smooth regions and a few sharp edges. Furthermore, because edges should have the same location in all MR contrasts, a multi-channel total-variation (MTV) distribution is used, which promotes modalities with common smoothness profiles [18]. Estimating the HR images given a set of observed LR images is then cast as an inference problem in this probabilistic model. In order for the model to generalise well, its hyper-parameters are either estimated from the observed data (Gaussian and MTV parameters) or set in a general and consistent way (slice-selection profile).

3.1 The Generative Model

Let $\mathbf{Y} = \{\{\mathbf{y}_m\}_{m=1}^{M} \mid \mathbf{y}_m \in \mathbb{R}^N\}$ denote the unknown HR images of M different MR contrasts and $\boldsymbol{\mathcal{X}} = \{\{\mathbf{X}_m\}_{m=1}^{M} \mid \mathbf{X}_m = \{\mathbf{x}_{ms}\}_{s=1}^{S_m} \mid \mathbf{x}_{ms} \in \mathbb{R}^{N_{ms}}\}$ a set of LR images, where S_m is the number of observed subject images of MR contrast m. The joint probability of the model can then be written as:

$$p(\boldsymbol{\mathcal{X}}, \mathbf{Y}) = p(\boldsymbol{\mathcal{X}}|\mathbf{Y})p(\mathbf{Y}) = \prod_{m=1}^{M} \prod_{s=1}^{S_m} p(\mathbf{x}_{ms}|\mathbf{y}_m)p(\mathbf{Y}). \tag{1}$$

4

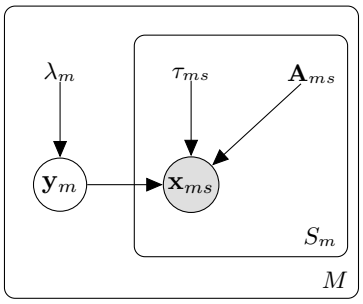

Fig. 1: Graphical model for multi-modality SR in MR scans using a MTV prior. Random variables are in circles. Observed variables are shaded. Plates indicate replication. There are M unknown HR images $\mathbf{y}_m$. For each HR image there are S_m observed LR images $\mathbf{x}_{ms}$. Model hyper-parameters are: Gaussian noise of precisions τ_{ms}, projection matrices $\mathbf{A}_{ms}$, and regularisation parameters λ_m.

A graphical representation of the joint probability in (1) is shown in Figure 1.

The conditional distribution of an observed LR image, given an unknown HR image, is assumed to be drawn from a likelihood function that is well-established in the SR literature, the multivariate Gaussian:

$$\begin{aligned} p(\mathbf{x}_{ms}|\mathbf{y}_m) &= \mathcal{N}(\mathbf{x}_{ms}|\mathbf{A}_{ms}\mathbf{y}_m, \tau_{ms}^{-1}\mathbf{I}) \\ &= \frac{\tau_{ms}^{N_{ms}/2}}{(2\pi)^{N_{ms}/2}} \exp\left(-\frac{\tau_{ms}}{2}(\mathbf{A}_{ms}\mathbf{y}_m - \mathbf{x}_{ms})^{\mathrm{T}}(\mathbf{A}_{ms}\mathbf{y}_m - \mathbf{x}_{ms})\right), \end{aligned} \tag{2}$$

where τ_{ms} is the precision of the observation noise ($\tau_{ms} = 1/\sigma_{ms}^2$) and $\mathbf{A}_{ms} \in \mathbb{R}^{N_{ms} \times N}$ a projection matrix. Furthermore, the reconstructions are assumed to all have the same dimensions and 1 mm isotropic voxel size.

The prior probability of the reconstructions is given by:

$$p(\mathbf{Y}) = Z^{-1} f(\mathbf{Y}), \tag{3}$$

and is defined by the normalising factor Z, which is independent of $\mathbf{Y}$, and the MTV prior, defined by:

$$f(\mathbf{Y}) = \prod_{n=1}^{N} \exp\left(-\sqrt{\sum_{m=1}^{M} \lambda_m \|\mathbf{D}_n \mathbf{y}_m\|_2^2}\right), \tag{4}$$

where $\mathbf{D}_n\mathbf{y}_m \in \mathbb{R}^3$ is the finite forward difference at voxel n, of reconstruction m. Using the MTV norm can intuitively be thought of as a method encouraging patterns where regions of larger gradient magnitude are similar over all MRI contrasts, through its sparsity inducing effect. Note that for $M = 1$, the MTV norm reduces to ordinary, isotropic TV.

3.2 Model Optimisation

A maximum a posteriori (MAP) estimate of the HR images can be obtained from the joint probability in (1) by minimising its negative logarithm:

$$\underset{\mathbf{Y}}{\operatorname{argmin}} \left\{ -\ln p(\mathbf{Y}|\boldsymbol{\mathcal{X}}) \right\}, \tag{5}$$

where

$$\begin{aligned} -\ln p(\mathbf{Y}|\boldsymbol{\mathcal{X}}) = & \sum_{m=1}^{M} \left(\sum_{s=1}^{S_m} \left(\frac{\tau_{ms}}{2} \left\| \mathbf{A}_{ms}\mathbf{y}_m - \mathbf{x}_{ms} \right\|_2^2 \right) \right) \\ & + \sum_{n=1}^{N} \left(\sqrt{\sum_{m=1}^{M} \lambda_m \left\| \mathbf{D}_n \mathbf{y}_m \right\|_2^2} \right) + \text{const.} \end{aligned} \tag{6}$$

The optimisation problem in (5) is solved by an alternating direction method of multipliers (ADMM) algorithm [19]. ADMM is part of a class of optimisation methods called proximal algorithms. ADMM recasts the energy minimisation as a constrained problem, from which an augmented Lagrangian can be formulated. The Lagrangian is then minimised in an alternating fashion until a convergence criterion is met. An ADMM algorithm was chosen because it is straightforward to implement and gives competitive results when solving TV problems in neuroimaging [20].

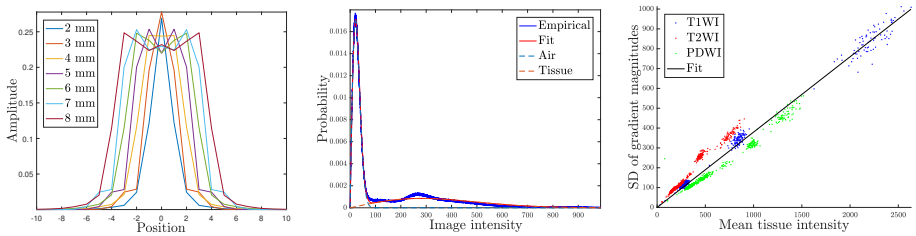

(a) **(b)** **(c)**

Fig. 2: Estimating the hyper-parameters of the proposed SR model. a) Example of modelled slice profiles for increasing slice thicknesses. b) Mixture of two Rician distributions fitted to the intensity histogram of an MR scan. c) Least-squares fit between mean tissue intensities and standard deviation of gradient magnitudes.

3.3 Setting of Hyper-Parameters

For a generative model to generalise well to new observations, required hyper-parameters need to be either modelled as random variables or estimated from the observed data. Without a principled way of choosing these parameter, the model

would not produce accurate results and a user would have to resort to manual tuning. Here, the hyper-parameters are the projection matrices, the noise precisions and the regularisation parameters. This section gives more detail about how these parameters are chosen.

Projection matrices The projection matrices $\mathbf{A}_{ms}$ incorporates into the model an assumption of how each LR image was generated. The slice-selection is what generates the thick-sliced LR images and is therefore what should be modelled by these matrices. Mathematically, the slice selection can be expressed as the result of a sequence of linear operators on an unknown HR image, consisting of a geometrical transformation, application of the slice-profile and down-sampling. That is, the sth LR image, from the mth modality, is assumed drawn from the following observation model:

$$\mathbf{x}_{ms} = \mathbf{A}_{ms}\mathbf{y}_m + \boldsymbol{\eta}_{ms} = \mathbf{D}_{ms}\mathbf{S}_{ms}\mathbf{y}_m + \boldsymbol{\eta}_{ms}, \tag{7}$$

where $\mathbf{S}$ simulates the slice-select profile of the MRI acquisition (see Figure 2a), $\mathbf{D}$ represents the down sampling of the HR image to the LR grid and $\boldsymbol{\eta}_{ms} \sim \mathcal{N}_{N_{ms}}(0, \tau_{ms}^{-1}\mathbf{I})$. Both $\mathbf{S}$ and $\mathbf{D}$ are dependent on a geometrical transformation $\mathbf{T} \in \mathbb{R}^{4\times 4}$ that can be read from the image header. The observation model in (7) is well established for SR of MR scans [8, 21].

Precision of the conditional distribution The trade-off between fitting the data and keeping the result smooth will depend on how noisy each LR image is. Because MR scans are reconstructed as the magnitude of an image that was originally complex, the assumed Gaussian noise model is just an approximation of the true Rician noise model. It is therefore of interest to estimate the amount of Rician noise in each observed MR scan, because this would in turn give estimates of the parameters τ_{ms}. In this work, a mixture of two Rician distributions is therefore fitted to the intensity histogram of each MR scan [22], and the precision of the observation noise is calculated from the class corresponding to air (see Figure 2b).

Regularisation parameter From a probabilistic modelling point of view, TV corresponds to the following Laplace distribution:

$$p(\mathbf{y}) = Z^{-1} f(\mathbf{y}) = Z^{-1} \prod_{n=1}^{N} \exp\left(-\lambda \left\|\mathbf{D}_n \mathbf{y}\right\|_2\right), \tag{8}$$

where $\lambda = 1/b$ and $b = \sqrt{\sigma^2/2}$. The hyper-parameter λ can therefore be related to the standard deviation (σ) of the gradient magnitude ($\|\mathbf{D}_n\mathbf{y}\|_2$) of the voxels in HR MR scans. This relationship is here used to obtain estimates of each λ_m in (4). More specifically, a least-squares fit is performed on the average tissue intensities, and the standard deviations of the gradient magnitudes, computed from a collection of 1,728 T1-, T2- and PD-weighted MR scans from the IXI dataset, resliced to have 1 mm isotropic voxel size (see Figure 2c). From this fit, a scaling function relating the image intensity of a LR image to the standard deviation of the gradient magnitude in a HR image can be constructed to give an estimate of a λ_m.

To summarise, a mapping for the regularisation parameters are learnt from HR images, the noise precisions are estimated from the LR data, and the projection matrices are based on metadata (slice thickness).

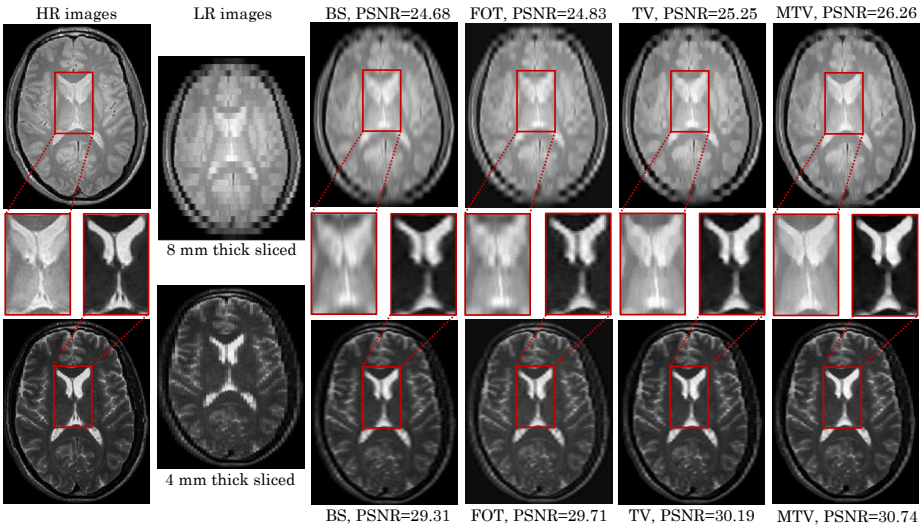

Fig. 3: Two LR images were simulated from two HR images (PD- and T2-weighted) and then reconstructed using four methods. BS, FOT and TV are single-modality SR techniques, capable of combining LR images of only one MR contrast. MTV on the other hand, uses information from both MR contrasts.

4 Experiments and Results

The proposed model was validated on the IXI dataset. LR images were generated from HR images allowing for ground-truth comparisons. The peak signal-to-noise ratio (PSNR), root-mean-square error (RMSE) and dice score were used to quantify the findings. Four methods for obtaining estimates of the unknown HR images were used. One interpolation based, using cubic b-splines (BS); three MAP based, using first-order Tikhonov (FOT), TV and MTV regularisation. FOT regularisation was achieved by changing the prior in (3) accordingly. Figure 3 shows an example of simulating LR images from HR images, and subsequently estimating the HR images from the LR ones using the different methods.

The following five experiments were conducted:

1. **Estimating the observation noise** 1,728 images were used to verify the fit of the Rician mixture model. For each image, a known percentage of Rician noise was added (1%, 2.5%, 5%, 10%). The Rician mixture model was then fitted to the intensity histogram of the noisy image. From the fit,

an estimate of the added noise was obtained. Table 1 shows the results of this estimation. From these results it is clear that the Rician mixture model accurately estimates a wide range of noise levels.

2. **Estimating the regularisation parameter** Four MR scans, from different subjects and of different contrasts (T1-, T2-, PD- and DTI-weighted), were used to validate whether the estimate of the regularisation parameter λ_m resulted in accurate reconstructions. The LR images were simulated having slice-thickness seven times greater than the in-plane resolution. Grid-searches over the regularisation parameters were then performed. Each grid-search was done in the range 10e{-4:0.2:1}. For each value of the regularisation parameter, the PSNR was computed between the resulting reconstruction and the known ground-truth. The DTI-weighted image was included to enable investigating if the method generalises to MR contrasts not part of the training set. The result of each grid-search, with the corresponding estimates marked by crosses, can be seen in Figure 4. This shows that the method for estimating the regularisation parameter produces good reconstructions, even for unseen MR contrasts.
3. **Increasing number of LR images** An experiment was conducted to verify that image quality improves for an increasing number of LR images. For 20 subjects, LR images were simulated with a slice thickness seven times greater than the in-plane resolution. Every three LR images were simulated having orthogonal thick slice directions. For more than three images, simulations additionally included a 1 mm translational shift. HR images were reconstructed using BS averaging, FOT and TV. Figure 5a shows the results of the experiment. It can be seen that the reconstruction quality increases as more LR images become available. Furthermore, TV consistently reconstructs the highest quality images. As for the optimal number of LR images, it seems as if four images would provide the best trade-off between acquisition time and reconstruction quality.
4. **Multi-modality SR** The subjects of the IXI dataset having T1-, T2- and PD-weigted images were used (576 subjects) to quantify the benefit of reconstructing from subjects imaged with multiple MR contrasts. For each subject's T1-, T2- and PD-weighted data, LR images were simulated by picking the thick-slice direction and the downsampling factor at random (varying from a factor of two to eight). HR images were then reconstructed using BS, FOT, TV and MTV. Table 2 shows the average PSNR where MTV reconstructed images with the greatest mean PSNRs and lowest standard deviation (sd). Figure 5b additionally shows average PSNR for different slice thicknesses, in which MTV once again performs favourably. However, for small slice thickness (2 mm) MTV was outperformed by both FOT and TV. This could be due to inaccurate estimates of the noise precisions and/or regularisation parameters, or misregistration of the LR images.
5. **Predicting age from brain segmentations** To assess the value of SR in segmentation tasks, all the reconstructions from Experiment 4, as well as the HR references, were segmented into grey matter, white matter and cerebrospinal fluid using the multi-channel segmentation algorithm in the SPM

software[4]. Dice scores were then calculated between the HR reference segmentation and the segmentations obtained from the reconstructions. The results can be seen in Table 3 where MTV obtained the highest dice score. All of the segmentations were also given as training data, together with age labels, to a support vector machine regressor from which age predictions were performed. Table 4 shows RMSEs in predicting age. Interestingly, the RMSE in age from training using the MTV reconstructed segmentations was actually smaller than the RMSE from using the HR reference data. This result could be due to smoothing introduced in the reconstructed data.

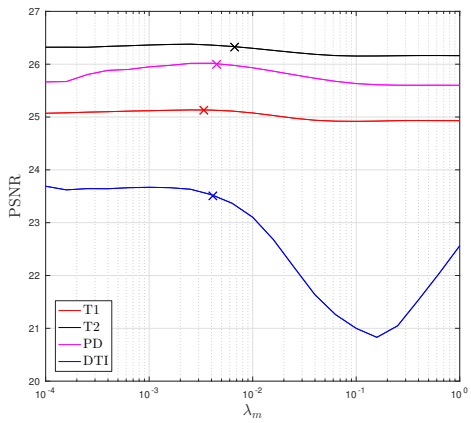

Fig. 4: Grid-search over the regularisation parameters λ_m, for four different MR contrasts. The estimates of the regularisation parameterers are marked by crosses.

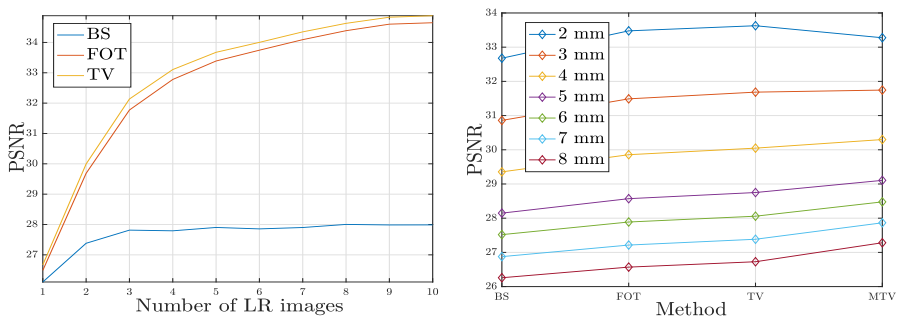

Fig. 5: a) PSNR for an increasing number of LR images. b) PSNR for increasing slice thicknesses.

[4] Available from http://www.fil.ion.ucl.ac.uk/spm/.

10

Table 1: Estimates of Gaussian noise precisions from 1,728 IXI images.

Ground-truth (%)	0	1	2.5	5	10
Estimate (%)	0.46 ± 0.49	1.12 ± 0.78	2.66 ± 0.71	5.00 ± 1.02	9.67 ± 3.52

Table 2: PSNR for different reconstruction methods. Shown as mean±std.

Modality	BS	FOT	TV	MTV
T1	29.83 ± 9.42	30.41 ± 10.70	30.62 ± 10.81	30.73 ± 9.35
T2	28.16 ± 8.15	28.63 ± 8.92	28.76 ± 8.94	29.02 ± 7.75
PD	28.46 ± 9.97	28.84 ± 10.07	29.02 ± 9.92	29.42 ± 8.46

Table 3: Dice scores for different reconstruction methods. Shown as mean±sd.

Class	BS	FOT	TV	MTV
GM	0.831 ± 0.002	0.839 ± 0.002	0.836 ± 0.002	0.848 ± 0.002
WM	0.867 ± 0.002	0.874 ± 0.001	0.871 ± 0.001	0.881 ± 0.001
CSF	0.852 ± 0.001	0.858 ± 0.001	0.855 ± 0.001	0.866 ± 0.001

5 Discussion

This paper presented a generative model for SR in routine clinical MRI. A validation on a large number of subjects demonstrated that the model can recover information from a collection of LR subject MR scans. The MTV norm was introduced as a prior for SR in MRI and was shown able to recover information between MR scans of different contrasts. The model can be used to recover resolution on any set of patient MR scans, of arbitrary MR contrast and orientation, with good results. This is because the necessary hyper-parameters are estimated from the observed data, eliminating any need for parameter tuning in order to obtain accurate HR reconstructions. As a lack of generalisably is a common drawback among published algorithms, this is a strength of the proposed model. Furthermore, as reconstruction results improve with the availability of more LR images (see Figure 5a), the model should generalise well to patient data, in which anatomical information is distributed over multiple LR scans.

The model could be of value in translating methods that have shown good results on research data to clinical imaging. For example, many techniques based on machine learning show promising results on analysing neuroimaging data. However, the amount of available high quality training images is limited, since it

Table 4: RMSE for predicting age using 558 multi-modality brain segmentations. Shown as mean±sd.

Predicted	HR	BS	FOT	TV	MTV
RMSE (years)	6.95 ± 0.96	7.00 ± 1.24	6.85 ± 1.27	6.71 ± 1.23	6.68 ± 1.21

is both time-consuming and expensive to acquire large quantities of HR scans. As hospitals contain huge amounts of population representative data, accumulated over years of clinical service and available for research at cost neutrality, resolution recovery on this data could provide these algorithms with an abundance of training material. Furthermore, the proposed model also has the potential of improving the results of multi-channel segmentation algorithms, in which interpolation schemes usually are necessary to reslice images to the same size. Here, the model has the potential of decreasing partial volume effects and to create more accurate resized images, which could improve the segmentation output.

Applying SR to clinical MRI is a complex problem and unavoidably leads to limitations in the proposed model. For example, MRI scanner parameters, such as slice profile and slice gap, are highly variable and assuming such parameters as fixed can lead to inaccuracies. However, it should be possible to instead estimate these parameters from the data, by including them in the optimisation procedure. Misalignment between scans could be dealt with in a similar manner, rather than by performing an initial rigid registration of the LR images, as is currently done. A well known fact in image SR is that TV introduces stair-casing effects on flat areas, which are abundant in brain MRI. Using other types of regularisation have shown effective in decreasing such artefacts [14]. However, the MTV regularisation proposed in this paper seems not to suffer from stair-casing artefacts. This is probably due to the fact that MTV penalises based on gradients distributed over MR scans with often, orthogonal thick-slice directions. Hence, an image area containing flat gradients in one MR contrast may very well have more informative gradients in another contrast. Finally, the model will need to be validated on real patient data, which should be done in cooperation with an expert physician. This validation will be part of future work.

References

1. C. Roobottom, G. Mitchell, and G. Morgan-Hughes, "Radiation-reduction strategies in cardiac computed tomographic angiography," *Clin Radiol*, vol. 65, no. 11, pp. 859–867, 2010.
2. S. M. Smith and T. E. Nichols, "Statistical challenges in big data human neuroimaging," *Neuron*, vol. 97, no. 2, pp. 263 – 268, 2018.
3. C. Blaiotta, P. Freund, M. J. Cardoso, *et al.*, "Generative diffeomorphic modelling of large MRI data sets for probabilistic template construction," *NeuroImage*, 2017.
4. M. Havaei, A. Davy, D. Warde-Farley, A. Biard, *et al.*, "Brain tumor segmentation with deep neural networks," *Med Image Anal*, vol. 35, pp. 18–31, 2017.
5. O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *MICCAI*, pp. 234–241, Springer, 2015.

12

6. E. Van Reeth, I. W. Tham, C. H. Tan, *et al.*, "Super-resolution in magnetic resonance imaging: A review," *Concepts Magn Reson*, vol. 40, no. 6, pp. 306–325, 2012.
7. E. Plenge, D. H. Poot, M. Bernsen, G. Kotek, *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*, vol. 68, no. 6, pp. 1983–1993, 2012.
8. H. Greenspan, G. Oz, N. Kiryati, *et al.*, "MRI inter-slice reconstruction using super-resolution," *Magn Reson Imag*, vol. 20, no. 5, pp. 437–446, 2002.
9. D. H. Poot, V. Van Meir, and J. Sijbers, "General and efficient super-resolution method for multi-slice MRI," in *MICCAI*, pp. 615–622, Springer, 2010.
10. J. V. Manjón, P. Coupé, A. Buades, *et al.*, "Non-local MRI upsampling," *Med Image Anal*, vol. 14, no. 6, pp. 784–792, 2010.
11. D. C. Alexander, D. Zikic, J. Zhang, *et al.*, "Image quality transfer via random forest regression: applications in diffusion MRI," in *MICCAI*, pp. 225–232, Springer, 2014.
12. S. Wang, Z. Su, L. Ying, *et al.*, "Accelerating magnetic resonance imaging via deep learning," in *ISBI*, pp. 514–517, IEEE, 2016.
13. M. Ebner, K. K. Chung, F. Prados, *et al.*, "Volumetric reconstruction from printed films: Enabling 30 year longitudinal analysis in MR neuroimaging," *NeuroImage*, vol. 165, pp. 238–250, 2018.
14. F. Odille, A. Bustin, B. Chen, *et al.*, "Motion-corrected, super-resolution reconstruction for high-resolution 3D cardiac cine MRI," in *MICCAI*, pp. 435–442, Springer, 2015.
15. J. V. Manjón, P. Coupé, A. Buades, *et al.*, "MRI superresolution using self-similarity and image priors," *Int J Biomed Imag*, vol. 2010, p. 17, 2010.
16. F. Rousseau, A. D. N. Initiative, *et al.*, "A non-local approach for image super-resolution using intermodality priors," *Med Image Anal*, vol. 14, no. 4, pp. 594–605, 2010.
17. K. Jafari-Khouzani, "MRI upsampling using feature-based nonlocal means approach," *IEEE Trans Med Imag*, vol. 33, no. 10, pp. 1969–1985, 2014.
18. Y.-W. Wen, M. K. Ng, and Y.-M. Huang, "Efficient total variation minimization methods for color image restoration," *IEEE Trans Image Process*, vol. 17, no. 11, pp. 2081–2088, 2008.
19. S. Boyd, N. Parikh, E. Chu, *et al.*, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found Trends Mach Learning*, vol. 3, no. 1, pp. 1–122, 2011.
20. E. D. Dohmatob, A. Gramfort, B. Thirion, *et al.*, "Benchmarking solvers for TV-1 least-squares and logistic regression in brain imaging," in *PRNI*, pp. 1–4, IEEE, 2014.
21. R. Z. Shilling, T. Q. Robbie, T. Bailloeul, *et al.*, "A super-resolution framework for 3-D high-resolution and high-contrast imaging using 2-D multislice MRI," *IEEE Trans Med Imag*, vol. 28, no. 5, pp. 633–644, 2009.
22. J. Ashburner and G. R. Ridgway, "Symmetric diffeomorphic modeling of longitudinal structural MRI," *Front Neurosci*, vol. 6, p. 197, 2013.