Published in final edited form as:

Spatiotemporal Image Anal Longitud Time Ser Image Data (2014). 2015 ; 8682: 67–76. doi: 10.1007/978-3-319-14905-9_6.

Longitudinal Guided Super-Resolution Reconstruction of Neonatal Brain MR Images

Feng Shi, Jian Cheng, Li Wang, Pew-Thian Yap, and Dinggang Shen

Department of Radiology and BRIC, University of North Carolina at Chapel Hill, NC, USA

Feng Shi: fengshi@med.unc.edu; Jian Cheng: dgshen@med.unc.edu

Abstract

Neonatal images have low spatial resolution and insufficient tissue contrast. Generally, interpolation methods are used to upsample neonatal images to a higher resolution for more effective image analysis. However, the resulting images are often blurry and are susceptible to partial volume effect. In this paper, we propose an algorithm that utilizes longitudinal prior information for effective super-resolution reconstruction of neonatal images. We use a non-local approach to learn the spatial relationships of brain structures in high-resolution longitudinal images and apply this information to the super-resolution reconstruction of the neonatal image. In other words, the recurring patterns throughout the longitudinal scans are leveraged for reconstructing the neonatal image with high resolution. To solve this otherwise ill-posed inverse problem, low-rank and total-variation regularizations are enforced. Experiments performed on both T1- and T2-weighted MR images of 28 neonates demonstrate that the proposed method is capable of recovering more structural details and outperforms methods such as nearest neighbor interpolation, spline-based interpolation, non-local means upsampling, and both low-rank and total variation based super-resolution.

1 Introduction

Spatial resolution of neonatal magnetic resonance (MR) images is limited by diverse factors such as imaging hardware, signal to noise ratio, and scanning time constraints [1]. High-resolution (HR) images with small voxel size are often desired for greater structural details [2]. In other words, images with low resolution (LR) are often affected by partial volume effect (PVE), where a voxel captures signal from multiple tissue types, resulting in fuzzy tissue boundaries [3]. This poses significant challenges for subsequent image analysis, for example, in the assessment of volumetric and shape changes of anatomical structures. PVE is especially severe in brain scans of neonates, due to their small brain size and intrinsically low tissue signal contrast.

Interpolation methods are commonly used to upsample neonatal images to a higher resolution before further analysis [4]. However note that each voxel in an LR image is

essentially a weighted average of corresponding voxels of a latent HR image. Thus, applying interpolation methods do not recover the HR image details with high frequency but causes further blurring to the image by performing another round of averaging on the voxels of the LR image. To address this issue, super-resolution (SR) techniques have been developed to estimate the HR image from one or more LR input images by reverting the image degradation process [1, 5]. Many existing approaches focus on single-frame SR, where only one LR image is available to recover the HR image. For example, non-local means upsampling was proposed for HR image reconstruction in [6]. In [7], both low-rank and total variation are used to regularize the otherwise ill-posed image reconstruction process. While these methods have been shown to be effective, using complementary information from multiple images might help improve reconstruction accuracy.

Longitudinal studies are widely employed to investigate the dynamic early brain structural and functional developments. In this setting, a subject is scanned for multiple times, such as at birth and 2 years of age. To address the challenges of low tissue contrast in neonatal images, recent studies have proposed to use their longitudinal follow-up images for guiding the image processing such as tissue segmentation [8]. The reason is that, the major brain myelination is established before birth while only fine-tuned after birth [9]. Fig. 1 shows a neonatal image and its 2-year-old image after affine alignment. Despite the differences in image contrast, brain structural patterns remain consistent longitudinally. Meanwhile, since the longitudinal images of a same subject share the identical brain anatomy, they could be better matched after registration than those images from different subjects.

In this paper, we propose a novel super-resolution method for recovering a HR neonatal image from a neonatal LR image using its longitudinal follow-up image as a prior. Specifically, since the follow-up images typically have higher resolution and tissue contrast, they are ideal for guiding the resolution enhancement of the neonatal brain images (Fig. 1). We first use a non-local approach to learn the spatial relationship of structures in high-resolution longitudinal images and then apply this information to the high-resolution reconstruction of the neonatal image. Our main contribution is three fold: 1) We learn longitudinal voxel relationship as a prior; 2) We integrate low-rank and total variation regularization for effective estimation of the HR image; 3) We explicitly model the image degradation processes involving blurring and downsampling. The proposed method will be evaluated using a group of neonatal images and compared with other state-of-the-art methods.

2 Method

We propose a novel method for neonatal image super-resolution reconstruction. First, we briefly introduce the super-resolution problem. Next, we put emphasis on the proposed non-local guidance from longitudinal prior. We then introduce the regularization terms as well as optimization steps for solving the cost function. The input will include a neonatal LR image and a longitudinal HR follow-up image, and the output will be the estimated neonatal HR image.

2.1 Super-Resolution Problem

We employ a physical model for capturing the degradation processes involved in reducing a high-resolution image to a low-resolution image [1]:

$$T = DSX + n \quad (1)$$

where T denotes the observed LR image, D is a downsampling operator, S is a blurring operator, X is the to-be-recovered HR image, and n represents the observation noise. The HR image can be estimated using this model by minimizing the following cost function:

$$\min_X \|DSX - T\|^2 + \lambda \mathfrak{R}(X) \quad (2)$$

where the first term is a data fidelity term used for penalizing the differences between the degraded HR image X and the observed LR image T . The second term is a regularization term often defined based on prior knowledge. Weight λ is introduced to balance the contributions of the fidelity term and regularization term.

2.2 Non-local Guidance from Longitudinal Prior

We use a non-local approach to learn the spatial relationships of structures in high-resolution longitudinal images and then apply this information to the reconstruction of high-resolution neonatal image. In other words, the recurring patterns throughout in the longitudinal scans are leveraged for reconstructing the neonatal image.

The non-local strategy has been proposed for image denoising [10]. For each voxel v in image X , the similarity between v and each voxel k in a non-local search domain $\Omega(v)$ is measured using their local patches. A weighted graph w can thus be obtained to represent the non-local relationships between v and other voxels in the large nonlocal search domain. Then, a non-local mean (NLM) image is obtained by updating each voxel in the image using this strategy:

$$NLM(X(v)) = \sum_{k \in \Omega(v)} w(v, k) X(k), \quad \forall v \quad (3)$$

In our case, X is the neonatal HR image that needs to be recovered. To utilize the longitudinal prior, we propose to learn the weighted graph w for each voxel using both the longitudinal HR image L and the pre-estimated X :

$$w(v, k) = \frac{1}{Z_v} e^{-\|P_L(v) - P_L(k)\|^2 / h_L^2} e^{-\|P_X(v) - P_X(k)\|^2 / h_X^2} \quad (4)$$

where $w(v, k)$ is the weight associating the center voxel v to a voxel k in its search domain $\Omega(v)$, $P_L(v)$ and $P_L(k)$ are 3D patches of longitudinal HR image L centered at v and k , $P_X(v)$ and $P_X(k)$ are 3D patches of image X centered at v and k , h_L and h_X are parameters controlling the strength of the weights, and Z_v is the normalization constant. We refer to the

obtained non-local mean image as $NLM(X, L)$. Note that the simultaneous consideration of both LR and HR images increases robustness to structural misalignments.

Fig. 2 illustrates the benefit of using the longitudinal prior. We first simulate a neonatal LR image by applying Gaussian blurring and downsampling to a real HR neonatal image, and then upsample it using spline-based interpolation. For a voxel near a gyrus, we compute its non-local similarity values in the interpolated neonatal LR image (Fig. 2A) and found that its non-local weight map is dominated by the gyrus structure in its neighborhood, which may be due to the blurring effect in the degraded image. On the other hand, the longitudinal HR image provides an informative pattern (Fig. 2B). After combining these two sets of weights, we obtain a new weight map (Fig. 2C) that resembles closely the results given by the original neonatal HR image (Fig. 2D). The computed weight map would be very useful for guiding the reconstruction process. Note that, an advantage of our method is the proposed nonlocal weights are calculated based on the distance of patches, and thus will be robust to the different contrast patterns in the neonatal and longitudinal images.

The non-local mean image is then used as a penalty term to guide the recovering of the neonatal HR image X :

$$\mathfrak{R}_{NLM}(X) = \lambda_{NLM} \|X - NLM(X, L)\|^2 \quad (5)$$

2.3 Global and Local Self-Similarity

We follow a previous study [7] to add regularization terms for ensuring both global and local self-similarity in the estimated HR neonatal image. The proposed method is thus reformulated as follows:

$$\min_X \|DSX - T\|^2 + \lambda_{NLM} \|X - NLM(X, L)\|^2 + \lambda_{rank} Rank(X) + \lambda_{tv} TV(X) \quad (6)$$

where the third term is for low-rank regularization, and the fourth term is for total variation regularization. λ_{rank} and λ_{tv} are the respective tuning parameters. These two additional terms are introduced below.

Low-Rank Regularization—Low-rank is an assumption often used in matrix completion tasks, where the matrix is incomplete and the goal is to estimate missing values from a small number of entries. Here we use low-rank as a regularization term to help retrieve useful information from remote regions. The rank of 3D image is defined as [11]:

$$Rank(X) = \sum_{i=1}^3 \alpha_i \|X_{(i)}\|_{tr}, \text{ where the rank is computed as the combination of trace norms of all matrices unfolded along each dimension. } \alpha_i \text{ are parameters satisfying } \alpha_i \geq 0 \text{ and } \sum_{i=1}^3 \alpha_i = 1. X_{(i)} \text{ is the unfolded } X \text{ along the } i\text{-th dimension.}$$

Total-Variation Regularization—Total-variation is defined as the integral of the absolute gradients of the image [12]: $TV(X) = \int |\nabla X| dx dy dz$. It is proven useful in image

denoising and super-resolution [12]. One of main advantages of total variation is its ability to effectively preserve edges in the image.

2.4 Optimization

We employ the alternating direction method of multipliers (ADMM) algorithm to optimize the cost function in Eq. (6). Following [13], we introduce variables $\{M_i\}_{i=1}^3$ and equality constraints $X_{(i)} = M_{i(i)}$, and thus the Lagrangian cost function is:

$$\min_{X, \{M_i\}_{i=1}^3, \{U_i\}_{i=1}^3} \|DSX - T\|^2 + \lambda_{NLM} \|X - NLM(X, L)\|^2 + \lambda_{rank} \sum_{i=1}^3 \alpha_i \|M_{i(i)}\|_{tr} + \sum_{i=1}^3 \frac{\rho}{2} (\|X - M_i + U_i\|^2 - \|U_i\|^2) \lambda_{tv} \int |\nabla X| dx dy dz \quad (7)$$

where $\{U_i\}_{i=1}^3$ are Lagrangian dual variables to integrate the equality constraints into the cost function. Then, we break the cost function into three subproblems and iteratively update them. Note that we propose to keep $NLM(X, L)$ fixed to achieve a convex solution when optimizing the three subproblems. The entire super-resolution process is summarized as Algorithm 1.

Subproblem 1—Update $X^{(k+1)}$ by minimizing:

$$\arg \min_X \|DSX - T\|^2 + \lambda_{NLM} \|X - NLM(X, L)\|^2 + \sum_{i=1}^3 \frac{\rho}{2} \|X - M_i^{(k)} + U_i^{(k)}\|^2 + \lambda_{tv} \int |\nabla X| dx dy dz \quad (8)$$

Subproblem 2—Update $\{M_i^{(k+1)}\}_{i=1}^3$ by minimizing:

$$\min_{\{M_i\}_{i=1}^3} \lambda_{rank} \sum_{i=1}^3 \alpha_i \|M_{i(i)}\|_{tr} + \sum_{i=1}^3 \frac{\rho}{2} \|X^{(k+1)} - M_i + U_i^{(k)}\|^2 \quad (9)$$

Subproblem 3—Update $\{U_i^{(k+1)}\}_{i=1}^3$ by:

$$U_i^{(k+1)} = U_i^{(k)} + (X^{(k+1)} - M_i^{(k+1)}) \quad (10)$$

Algorithm 1

Longitudinal Guided Super-Resolution Reconstruction of Neonatal Images

Input: Low-resolution neonatal image T , high-resolution longitudinal image L

Initialize the desired high-resolution neonatal image X by upsampling T with spline-based interpolation. Set redundant variables $M_i = 0$, $U_i = 0$, $i = 1, 2, 3$.

Repeat

Update $NLM(X, L)$ using non-local weights computed from X and L ;

Repeat

Update X based on Eq. (8) using gradient descent;

Update M_i based on Eq. (9) using Singular Value Thresholding (SVT) [14];

Update U_i based on Eq. (10);

End

Until iteration difference in the cost function (Eq. (7)) is less than ε ;

End

Output: Reconstructed high-resolution neonatal image X ;

3 Experiments

3.1 Data

A total of 28 healthy infants (11 males and 17 females) were used in this study. They were firstly scanned at birth, and a follow-up scan was performed at 2 years of age. A Siemens head-only 3T scanner was used with a circular polarized head coil. T2 images were acquired with 58 axial slices at the resolution of $1.25 \times 1.25 \times 1.95 \text{ mm}^3$. T1 images were also acquired with 144 sagittal slices at the resolution of $1 \times 1 \times 1 \text{ mm}^3$.

All images were preprocessed using a standard image-processing pipeline, including bias correction and skull stripping [15]. T2 images were linearly aligned to their corresponding T1 images. The longitudinal follow-up images were also aligned to their neonatal images using affine registration followed by nonlinear diffeomorphic demons registration [16].

3.2 Experimental Setting

For evaluation of the proposed method, we simulated a group of neonatal LR images (Fig. 3) by applying blurring and downsampling operators to the original neonatal images. Images reconstructed by the proposed method were compared with the respective original images serving as ground-truth. Specifically, blurring was performed using a Gaussian kernel with standard deviation of 1 voxel. Downsampling was carried out by averaging every 8 voxels in an image, to simulate the partial volume effect. Signal-to-noise ratio (SNR) was used to compare the recovered image with the original image to evaluate the quality of reconstruction: $SNR = 20 * \log_{10}(\|Truth\|/\|Truth - Recovered\|)$. Higher SNR means better reconstruction performance.

Parameters were defined experimentally. We set $\alpha_1 = \alpha_2 = \alpha_3 = 1/3$, $\lambda_{rank} = 0.01$, $\lambda_{TV} = 0.01$, $\lambda_{NLM} = 0.02$, $dt = 0.1$, h_L and h_X were set at 1% intensity range with respect to images L and X , and the maximum iteration number was 200. The 3D patch size was $3 \times 3 \times 3$ voxels, and search domain was $7 \times 7 \times 7$ voxels. The difference between iterations was measured and the program stopped when the difference was less than $\varepsilon = 1e - 5$.

3.3 Results

Experiments were performed using two imaging modalities, i.e., T1 and T2, respectively. In each experiment, the HR neonatal images were reconstructed with the help of their corresponding longitudinal images of same modality. Methods for comparison include nearest neighbor interpolation (NN), spline interpolation (Spline), non-local means upsampling (NLM) [6], super-resolution method regularized by low-rank and total variation (LRTV) [7], and the proposed method. The implementation of NLM provided on author's

website was used¹. We implemented LRTV by setting $\lambda_{NLM} = 0$ in the proposed method. Spline-based interpolation was used as initialization for NLM, LRTV, and the proposed method.

Fig. 4 demonstrates representative reconstruction results for T1 (top panel) and T2 (bottom panel). For each modality, the left panel shows the input neonatal LR image and its longitudinal HR image at 2 years of age. The right panel shows the results from all methods. Close-up views of selected regions are also shown for better visualization. It can be observed that the results of NN and spline interpolation methods show severe blurring artifacts. NLM results also appear blurry, which is partly because the blurring degradation is not explicitly considered [6]. LRTV and the proposed method demonstrate edge-preserved results. The proposed method achieves the highest SNR.

Quantitative results for T1 and T2 reconstructions of 28 infants are shown in Fig. 5. Similar to the observations from Fig. 4, Fig. 5 indicates that SNR increases from NN, spline, NLM, LRTV, to the proposed method. The proposed method achieved the highest SNR among all methods ($p < 0.01$). Note that T1 images give relatively lower SNR than T2 images, which may be due to the fact that T1 images generally have lower tissue contrast in neonates. We also tried other strategies such as using longitudinal T2 images as prior for neonatal T1 image reconstruction, or using longitudinal T1 images as prior for neonatal T2 image reconstruction. These results are similar to those obtained with the use of priors from the same modality ($p > 0.05$).

4 Conclusion

We have presented a super-resolution method that combines both self-similar information within an image and the longitudinal information from its follow-up image. High-resolution reconstruction results of neonatal images indicate that the proposed method outperforms methods such as NN interpolation, spline-based interpolation, NLM upsampling, and LRTV-based super-resolution. Currently, the proposed method only works on images with longitudinal follow-ups. In future, we anticipate using the results as training data to learn their longitudinal developmental constraints and thus extend the method on single time-point images.

References

1. Park SC, Park MK, Kang MG. Super-resolution image reconstruction: a technical overview. *IEEE Signal Processing Magazine*. 2003; 20:21–36.
2. Freeman WT, Jones TR, Pasztor EC. Example-based super-resolution. *IEEE Computer Graphics and Applications*. 2002; 22:56–65.
3. Pham DL, Xu C, Prince JL. Current methods in medical image segmentation. *Annu Rev Biomed Eng*. 2000; 2:315–337. [PubMed: 11701515]
4. Lehmann TM, Gonner C, Spitzer K. Survey: interpolation methods in medical image processing. *IEEE Trans Med Imaging*. 1999; 18:1049–1075. [PubMed: 10661324]

¹<https://sites.google.com/site/pierrickcoupe/>

5. Farsiu S, Robinson D, Elad M, Milanfar P. Advances and challenges in super-resolution. *International Journal of Imaging Systems and Technology*. 2004; 14:47–57.
6. Manjón JV, Coupé P, Buades A, Fonov V, Louis Collins D, Robles M. Non-local MRI upsampling. *Medical Image Analysis*. 2010; 14:784–792. [PubMed: 20566298]
7. Shi, F.; Cheng, J.; Wang, L.; Yap, P-T.; Shen, D. Low-Rank Total Variation for Image Super-Resolution. In: Mori, K.; Sakuma, I.; Sato, Y.; Barillot, C.; Navab, N., editors. MICCAI 2013, Part I. LNCS. Vol. 8149. Springer; Heidelberg: 2013. p. 155-162.
8. Shi F, Fan Y, Tang S, Gilmore JH, Lin W, Shen D. Neonatal brain image segmentation in longitudinal MRI studies. *Neuroimage*. 2010; 49:391–400. [PubMed: 19660558]
9. Armstrong E, Schleicher A, Omran H, Curtis M, Zilles K. The Ontogeny of Human Gyrification. *Cerebral Cortex*. 1995; 5:56–63. [PubMed: 7719130]
10. Buades A, Coll B, Morel J-M. A non-local algorithm for image denoising. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*. 2005; 2:60–65.
11. Liu J, Musialski P, Wonka P, Ye J. Tensor Completion for Estimating Missing Values in Visual Data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2013; 35:208–220. [PubMed: 22271823]
12. Rudin LI, Osher S, Fatemi E. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*. 1992; 60:259–268.
13. Boyd S, Parikh N, Chu E, Peleato B, Eckstein J. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*. 2011; 3:1–122.
14. Cai JF, Candès EJ, Shen Z. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization*. 2010; 20:1956–1982.
15. Shi F, Wang L, Dai Y, Gilmore JH, Lin W, Shen D. LABEL: Pediatric Brain Extraction using Learning-based Meta-algorithm. *Neuroimage*. 2012; 62:1975–1986. [PubMed: 22634859]
16. Vercauteren T, Pennec X, Perchant A, Ayache N. Diffeomorphic demons: Efficient non-parametric image registration. *NeuroImage*. 2009; 45:S61–S72. [PubMed: 19041946]

Fig. 1.

T1 MR images of a neonate (left) and its follow-up at 2 years of age (right). The 2-year-old image was registered to the neonatal image using affine alignment. Two brain regions marked with green and red were zoomed up for close comparison.

Fig. 2.

Illustration of the proposed non-local weights. (A) shows a simulated neonatal LR image after degradation process. For the voxel marked by red, its local region is marked by green and shown in close-up view. The non-local weights of that voxel with voxels in search domain of $7 \times 7 \times 7$ are shown. The used image patch size is $3 \times 3 \times 3$ for each voxel. (B) is its longitudinal follow-up image. (C) is the combined weights from (A) and (B). (D) is the original neonatal HR image.

Fig. 3.

Simulation of a low-resolution image from an original neonatal image. The observed image will be the input for super-resolution reconstruction and the recovered image will be compared with original neonatal image for performance evaluation.

Fig. 4.

Reconstruction results for a neonatal T1 image (top row) and a neonatal T2 image (bottom row) from different subjects. In each row, the left panel shows input images including a low-resolution neonatal image and its 2-year-old follow-up image. The low-resolution neonatal image was simulated from real images (Fig. 3) showing at right panel tagged as Truth. The reconstructed high-resolution images from different methods were compared with the Truth image and their SNR were provided measuring reconstruction performance.

Fig. 5. SNR boxplots for reconstruction of 28 neonatal images using (A) T1 images and (B) T2 images. The proposed method significantly outperforms all other methods ($p < 0.01$).