NeuroImage 84 (2014) 939–950

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

Fiber-driven resolution enhancement of diffusion-weighted images

Pew-Thian Yap *, Hongyu An, Yasheng Chen, Dinggang Shen

Department of Radiology and Biomedical Research Imaging Center (BRIC), University of North Carolina at Chapel Hill, NC 27599, USA

ARTICLE INFO

Article history:
Accepted 10 September 2013
Available online 21 September 2013

Keywords:
Diffusion magnetic resonance imaging (DMRI)
Resolution enhancement
Anisotropic interpolation

ABSTRACT

Diffusion-weighted imaging (DWI), while giving rich information about brain circuitry, is often limited by insufficient spatial resolution and low signal-to-noise ratio (SNR). This paper describes an algorithm that will increase the resolution of DW images beyond the scan resolution, allowing for a closer investigation of fiber structures and more accurate assessment of brain connectivity. The algorithm is capable of generating a dense vector-valued field, consisting of diffusion data associated with the full set of diffusion-sensitizing gradients. The fundamental premise is that, to best preserve information, interpolation should always be performed along axonal fibers. To achieve this, at each spatial location, we probe neighboring voxels in various directions to gather diffusion information for data interpolation. Based on the fiber orientation distribution function (ODF), directions that are more likely to be traversed by fibers will be given greater weights during interpolation and vice versa. This ensures that data interpolation is only contributed by diffusion data coming from fibers that are aligned with a specific direction. This approach respects local fiber structures and prevents blurring resulting from averaging of data from significantly misaligned fibers. Evaluations suggest that this algorithm yields results with significantly less blocking artifacts, greater smoothness in anatomical structures, and markedly improved structural visibility.

Introduction

Diffusion-weighted imaging (DWI) (Johansen-Berg and Behrens, 2009) is a key imaging technique for the investigation and characterization of white matter pathways in the brain. It probes water diffusion in various directions and at various diffusion scales to characterize microstructural compartments that are much smaller than the voxel size. However, limited by today's imaging technique, the typical $(2\ \text{mm})^3$ resolution achievable by DMRI is too coarse to sufficiently capture the subtlety of neuronal axons, diameters of which range from 1 μm to 30 μm (Johansen-Berg and Behrens, 2009; Scherrer et al., 2012; Yap and Shen, 2012a). This causes significant partial volume effect since the signal collected at each voxel is likely to be due to multiple fascicles that concurrently traverse the voxel. Acquiring images with resolution higher than the typical $(2\ \text{mm})^3$, however, is extremely difficult without incurring unrealistic scan times and causing very low SNR due to reduced voxel size (Scherrer et al., 2012). The impact of noise is aggravated in high angular resolution diffusion imaging (HARDI), which often requires prolonged echo time (TE) to achieve relatively high diffusion weighting.

Increasing the resolution is not only important for registration, segmentation, and tractography to be performed with greater accuracy, but is also crucial for better visualization of anatomical structures to identify possible neuropathologies. Solutions to achieve higher resolution include employing higher magnetic fields or stronger/faster gradients, dedicated acquisition techniques (Heidemann et al., 2012; Liu et al., 2004; Scherrer et al., 2012; Sotiropoulos et al., 2013), as well as post-processing algorithms (Calamante et al., 2010; Gupta et al., 2013; Manjón et al., 2010a; Nedjati-Gilani et al., 2008). In the current work, we will take the post-processing approach, since this approach does not rely on expensive scanner upgrades and complex time-consuming sequences and can hence be applied to existing data without requiring re-acquisition.

In this paper, we present a technique that will exploit the continuity information given by local fiber architectures to increase the spatial resolution of diffusion-weighted data beyond the acquisition resolution. Our algorithm will generate a spatially dense vector-valued field consisting of diffusion data associated with the full set of diffusion-sensitizing gradient directions. Similar to (Calamante et al., 2010), our approach gains spatial resolution by using additional information obtained from outside of each individual voxel. Dissimilar to (Gupta et al., 2013; Nedjati-Gilani et al., 2008), our approach does not assume any diffusion model and is applied directly to the DW images. Our approach does not require special acquisition techniques, such as those proposed in (Greenspan, 2009; Greenspan et al., 2002; Scherrer et al., 2012), does not need expensive scanner hardware upgrades (to 7 T or 11 T), and can be applied to existing data. The key highlights of our method are as follows:

1. Directional Profiling — Our approach uses a directional profiling scheme to examine the neighborhood of each spatial location and to gauge the probability of whether a specific direction is likely to be traversed by fibers. The resulting directional probability

* Corresponding author.
E-mail address: ptyap@med.unc.edu (P.-T. Yap).

1053-8119/$ – see front matter © 2013 Elsevier Inc. All rights reserved.
http://dx.doi.org/10.1016/j.neuroimage.2013.09.016

distribution is then employed to encourage interpolation along tangential and not orthogonal directions of axonal fibers. Unlike the conventional trilinear interpolation, which does not take into account the directional nature of DWI data, our approach mimics DWI acquisition mechanism more closely by borrowing information from different directions to reconstruct the DWI data at each spatial location.

2. Microstructure-Preserving Smoothing — DWI data are typically noisy and need to be smoothed for increasing SNR. When smoothing, it is important that the boundaries defining the spatial extent of individual structures are preserved. In contrast to many existing methods that use *inter*-voxel gradient information to constrain smoothing to relatively homogeneous regions, our method uses *intra*-voxel fiber orientation distribution functions (ODFs) to guide smoothing. By constraining interpolation along fiber streamlines, we preserve not only the boundaries of the white matter, but also those between fiber tracts within the white matter.
3. Complete DWI Data — Our approach generates a vector-valued field of DWI data corresponding to the full set of diffusion-gradients and is not limited to the white matter. Gray matter provides contextual information for the white matter and the availability of gray matter data allows tissue segmentation based on diffusion data, such as that done in (Liu et al., 2007), to be performed, providing complementary tissue contrast to tissue segmentation based on structural MRI. The fact that our method produces a complete set of DWI data also allows any diffusion models to be fitted to the resolution-enhanced data for multifaceted analysis.

A preliminary version of this work was presented in our conference paper (Yap and Shen, 2012a). Herein, we provide additional examples, results, derivations, and insights that are not part of the conference publication.

Paper organization

In the upcoming sections, we will first detail in the Materials and methods section the key components of the proposed algorithm. We will then demonstrate the effectiveness of the proposed algorithm in the Results section with both in silico and in vivo data. Additional discussion is provided in the Discussion section before the paper is concluded in the Conclusions section.

Material and methods

Datasets

Various datasets were acquired or generated for comprehensive evaluation of the proposed method. They are described as follows.

In silico data

To quantitatively evaluate the accuracy of the proposed method, we generated a 96×96 field of diffusion-weighted signal, forming a spiral as shown in Fig. 2(A). Each voxel within the spiral was simulated using a tensor model with principal diffusivities $\lambda_1 = 1.5 \times 10^{-3}$ mm^2/s, $\lambda_2 = \lambda_3 = 3 \times 10^{-4}$ mm^2/s, and diffusion weighting $b = 2000$ s/mm^2. The baseline non-diffusion-weighted signal S_0 was set to 150. Diffusion-weighted signal was sampled along each of the 120 gradient directions obtained from the in vivo dataset (see the In vivo data section). The background voxels that fall outside the spiral were generated via isotropic diffusion with constant signal magnitude $S_0 \exp(-b\lambda)$, where $b = 2000$ s/mm^2, $\lambda = 2.5 \times 10^{-3}$ mm^2/s, and $S_0 = 1000$. Note that these diffusion parameters were carefully chosen to mimic the in vivo dataset described in the In vivo data section.

A cross phantom, shown in Fig. 2(C), was also used to evaluate the proposed method in situations where fiber crossings exist. Using the same diffusion parameters described above for the spiral, we used diffusion tensors and their mixtures to generate a phantom of size 48×48 for evaluation. One group of tensors was oriented in the horizontal direction and another group at an angle, i.e., 30°, 40°,..., 90°, with the horizontal direction. At locations where these two groups cross, a mixture of two tensors with equal volume fraction was used to model the crossings.

In vivo data

Diffusion-weighted images for 4 adult subjects were acquired using a Siemens 3 T TIM Trio MR scanner with an EPI sequence. Diffusion gradients were applied in 120 non-collinear directions with diffusion weighting $b = 2000$ s/mm^2, repetition time (TR) = 12,400 ms, and echo time (TE) = 116 ms. The imaging matrix was 128×128 with a field of view (FOV) of 256×256 mm^2. The slice thickness was 2 mm. Six non-diffusion-sensitized images ($b = 0$ s/mm^2) were acquired. T_1-weighted structural images with 1 mm isotropic resolution were also acquired as anatomical references.

High-resolution in vivo data

For further evaluation, a set of high-resolution $(1\text{ mm})^3$ diffusion-weighted images was acquired using the Siemens 3 T TIM Trio MR scanner with the acquisition technique reported in (Porter and Heidemann, 2009). Diffusion gradients were applied in 42 non-collinear directions with diffusion weighting $b = 1000$ s/mm^2. The imaging matrix was 192×192 with a field of view of 192×192 mm^2. The slice thickness was 1 mm.

Neonatal data

Diffusion-weighted images of a neonate were acquired at approximately one month after birth. Diffusion gradients were applied in 42 non-collinear directions with diffusion weighting $b = 1000$ s/mm^2, TR = 7680 ms and TE = 82 ms. The scans covered the whole brain with a resolution of $2 \times 2 \times 2$ mm^3.

Fiber-driven resolution enhancement

To increase spatial resolution, the image domain is uniformly divided using a grid with grid elements that are smaller than the acquisition voxel size. The diffusion-weighted data for each of these grid elements are then generated using the following steps: 1) Directional profiling in a field of fiber ODFs; 2) Interpolation of diffusion-weighted data based on the fiber orientation profile (generated in the previous step) with bias correction (owing to the Rician distribution nature of the magnitude signal); and 3) Mean-shift refinement for recovering more structural details. Each step is detailed in the following sections.

Local fiber profiling

Interpolation along directions transversed by fibers preserves structural boundaries. To determine the probability of whether a grid element at spatial location $\mathbf{x}$ is traversed by fibers in directions $\mathbf{v}_k$ ($k = 1, \ldots, M$), which are densely distributed on the unit sphere and not necessarily antipodal symmetric, we profile for each k the field of fiber ODFs $\{p(\mathbf{x}_i, \mathbf{v}) | \mathbf{x}_i \in \mathcal{N}(\mathbf{x})\}$ along direction $\mathbf{v} = \mathbf{v}_k$ (see Fig. 1), where $\mathcal{N}(\mathbf{x})$ is a neighborhood of voxels in the vicinity of $\mathbf{x}$. Note that $\mathbf{x}_i$ is a point in space at which the diffusion-weighted signal is actually acquired, and $\mathbf{x}$ is a point corresponding to a grid element of the high-resolution grid, using which the resolution-enhanced data will be reconstructed. The local fiber configuration at $\mathbf{x}$ is characterized by a *fiber orientation profile*, which is a directional function that allows for anisotropic interpolation of neighboring information to generate the DWI data of the grid element at location $\mathbf{x}$. It is determined as

$$\hat{p}(\mathbf{x}, \mathbf{v}_k) = \frac{\sum_{\mathbf{x}_i \in \mathcal{N}(\mathbf{x})} w(\mathbf{x}_i, \mathbf{x}, \mathbf{v}_k) p(\mathbf{x}_i, \mathbf{v}_k)}{\sum_{\mathbf{x}_i \in \mathcal{N}(\mathbf{x})} w(\mathbf{x}_i, \mathbf{x}, \mathbf{v}_k)} = \sum_{\mathbf{x}_i \in \mathcal{N}(\mathbf{x})} \tilde{w}(\mathbf{x}_i, \mathbf{x}, \mathbf{v}_k) p(\mathbf{x}_i, \mathbf{v}_k), \quad (1)$$