Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Magnetic Resonance 177 (2005) 102–110
JMR Journal of Magnetic Resonance
www.elsevier.com/locate/jmr

Investigations on the efficiency of cardiac-gated methods for the acquisition of diffusion-weighted images

Rita G. Nunes, Peter Jezzard, Stuart Clare *

FMRIB Centre, Department of Clinical Neurology, University of Oxford, John Radcliffe Hospital, Oxford OX3 9DU, UK

Received 14 April 2005; revised 30 June 2005
Available online 19 August 2005

Abstract

Diffusion-weighted images are inherently very sensitive to motion. Pulsatile motion of the brain can give rise to artifactual signal attenuation leading to over-estimation of the apparent diffusion coefficients, even with snapshot echo planar imaging. Such miscalculations can result in erroneous estimates of the principal diffusion directions. Cardiac gating can be performed to confine acquisition to the quiet portion of the cycle. Although effective, this approach leads to significantly longer acquisition times. On the other hand, it has been demonstrated that pulsatile motion is not significant in regions above the corpus callosum. To reduce acquisition times and improve the efficiency of whole brain cardiac-gated acquisitions, the upper slices of the brain can be imaged during systole, reserving diastole for those slices most affected by pulsatile motion. The merits and disadvantages of this optimized approach are investigated here, in comparison to a more standard gating method and to the non-gated approach.

Keywords: Diffusion-weighted imaging; Artifacts; Cardiac gating; Efficiency; Brain

1. Introduction

To obtain information regarding the structure of white matter, diffusion-weighted images are made sensitive to the microscopic movement of water molecules as they diffuse through tissue. This means, however, that the images also become extremely sensitive to any other form of movement that may be present. Image artifacts can therefore occur as a result both of bulk motion of the subject [1] or the pulsatile motion of the brain itself [2,3].

Pulsatile brain motion occurs as a result of the arterial expansion which follows systole [4] and is nonlinear in nature as different areas of the brain have different velocity profiles. Phase-velocity measurements have shown that the highest velocities are present in the inferior and medial areas of the brain and decrease towards the periphery where the tissue remains practically motionless throughout the cycle [4,5]. Deep brain regions are displaced in a centripetal and caudal (head to feet) way and the magnitude of the deformations can be very significant when compared to the diffusion distance, with displacements reaching as much as 0.1–0.13 mm and peak velocities 1–1.5 mm/s [4].

Snapshot echo planar imaging (EPI) is the standard technique used to acquire diffusion-weighted (DW) images, as the acquisition is sufficiently fast for bulk patient motion to become negligible. The time scale of the nonlinear component of motion is, however, in the order of tens to hundreds of milliseconds [5] and therefore comparable to the period necessary to apply diffusion sensitization. Artifacts can consequently still be visible even in EPI images [2,6,7]. Velocity gradients within the image voxels give rise to phase dispersion resulting in additional signal loss and over-estimation of the apparent diffusion coefficients [5].

To estimate the direction of main diffusivity it is necessary to acquire multiple DW images sensitized to

* Corresponding author. Fax: +44 1865 222717.
E-mail address: stuart@fmrib.ox.ac.uk (S. Clare).

1090-7807/$ - see front matter © 2005 Elsevier Inc. All rights reserved.
doi:10.1016/j.jmr.2005.07.005

diffusion along a set of different directions. If no gating is employed, a fraction of the data will be collected during systole. Under these circumstances about 20% of the images are likely to display severe artifacts [6], but an additional fraction of the images may also be affected to a lesser degree [6].

If a certain region of the brain affected by pulsatile motion is imaged during systole, estimates of diffusion will tend to be greater along the direction of motion, biasing the estimate of the principal diffusion direction. In lower regions of the brain, such as the cerebellum and the brainstem, the motion occurs mainly along the superior–inferior orientation [4]. As such, whenever the diffusion gradient direction contains a significant component along the SI direction the resulting data are likely to be compromised (diffusivity over-estimated) if acquisition occurs during systole. This will in turn result in a bias in the estimate of the direction of principal diffusivity. An example of this phenomenon can be seen in the higher content of blue in the color-coded maps presented by Pierpaoli et al. [2] for images of the cerebellum acquired during systole. Biasing may also be observed along directions other than the superior–inferior, as reported by Jones and Pierpaoli [8]. For instance, in the genu of the corpus callosum a strong component of motion occurs along the anterior–posterior direction [4]. In this region, the estimate of the direction of principal diffusivity may be biased towards the anterior–posterior direction instead of laterally when data are acquired during systole. If fiber-tracking is performed under these circumstances, the connection between the two hemispheres may even appear to have been disrupted. The presence of artifacts may therefore have important implications when performing tractography [8].

One possible solution to minimize artifacts due to pulsatile motion consists of triggering the acquisition, such that the images are only acquired during the quiet portions of the cardiac cycle [6,9]. Although effective, this strategy lengthens the acquisition time considerably, and this is why non-gated approaches are still preferred by many researchers.

In previous studies which investigated the nature and extent of pulsatile brain motion, no movement was detected above the ventricles [4,5]. This is in agreement with the observations made by Skare and Andersson [6]. They reported no significant differences in the standard deviation of the signal in areas located above the corpus callosum when comparing multiple repeats of the same image obtained both with and without cardiac gating [6].

As previous studies have already investigated the impact of pulsatile brain motion on both tensor derived quantities [2] and on fiber tracking [8], demonstrating how cardiac gating may help to reduce these effects, the focus of this paper was to investigate whether it is possible to optimize cardiac-gating acquisitions so as to reduce the total acquisition times. Since regions of the brain located above the corpus callosum remain practically motionless during the cardiac cycle they can be imaged during systole without significant artifacts being observed. By selecting the order in which the slices are acquired so that these upper areas are imaged during the critical portion of the cycle, a larger number of images can be collected per cycle. We demonstrate that this approach can lead to a significant improvement in scanning efficiency compared to the standard gating method [6]. Using this method only a small increase in acquisition time is required for the acquisition of images free from pulsatile motion artifacts, compared to the non-gated approach.

2. Methods

Data were acquired on a 3.0 T Varian Inova scanner. A birdcage radiofrequency head coil was used for both pulse transmission and signal detection. A diffusion-weighted snapshot EPI sequence was employed. To minimize eddy currents, a doubly refocused spin-echo sequence was chosen [10]. The diffusion gradients achieved a maximum strength of 21.3 mT/m for a b-value of 1000 s/mm^2. The slice thickness was 2.5 mm. The other acquisition parameters were: TE = 106 ms, bandwidth of 125 kHz, field of view of 240×240 mm^2, matrix size of 96×62 (half K-space acquisition reconstructed using hermitian-conjugation and subsequently interpolated to 128×128) corresponding to an effective in-plane resolution of 2.5×2.5 mm^2.

Peripheral gating was preferred as it is easier to use. Using peripheral gating it is necessary to take into account the time difference that exists between the detection of the peripheral R wave and the period for which a maximum level of artifacts is observed. As this delay may depend on the actual setup being used, a preliminary experiment was performed following the approach described in [6]. Diffusion-weighted images encompassing the region of the corpus callosum were acquired for a set of delays: 0, 200, 400, 600, and 800 ms (one slice per cardiac cycle). The images were sensitized to diffusion along the inferior–superior direction and multiple repeats (25 in our case) were collected to evaluate the variance of the signal. For six subjects with heart rates in the range 50–70 bpm, no artifacts were found for delays lower or equal to 400 ms while increased signal variance was detected for delays of 600 and 800 ms (data not shown).

As for the main experiment, each data set consisted of 40 DW images with the diffusion gradients applied along the z (through slice) direction. This direction was chosen as it corresponds to the highest peak velocities observed during systole [4,5]. The images sensitized to diffusion along this orientation should therefore be more prone

to artifacts. Although such a limited choice of diffusion orientations is not normally used in practice, as fiber-tracking cannot be performed using a single direction, having an increased sensitivity to pulsatile motion enables a harsher test when assessing the quality of the data.

The repetition time for the non-gated data was set to the minimum possible (TR = 8.0 s for 42 slices). An extra volume with no diffusion-weighting was acquired for registration purposes.

To be able to register the images into a common space, a 3D T_1-weighted FLASH image was acquired having a 1 mm^2 in-plane resolution and either 1 or 1.5 mm slice thickness. Parameters for the acquisition were: field of view of 256×192 mm^2, matrix size of 256×192, TE = 5 ms, TR = 15 ms, flip angle of 15°, 128 slices.

Three data sets were acquired for each subject: non-gated, gated using the optimized method and gated using the standard method [6]. The full protocol was used on seven normal volunteers (three males and four females).

In both gating schemes, gating was performed such that triggering occurred in every cycle.

The total acquisition time per slice on our system is 190 ms, accounting also for extra time in between acquisitions to avoid gradient coil heating. To avoid systole while using the standard method, three slices were acquired per cardiac cycle.

For the optimized scan, the order in which the slices were acquired was adjusted to ensure that the upper slices of the brain would be scanned during systole. The number of slices scanned per cycle was set depending on the maximum estimated heart rate, so as to minimize the acquisition time. The total number of slices could also be increased from a minimum of 42 so as to be a multiple of the number of slices acquired per cycle. For example, for a maximum estimated heart rate of 60 bpm, five slices can be acquired per heart beat and the total number of slices increased to 45. These are divided into 5 groups of 9 slices each as shown in Fig. 1 (labeled from 'A' to 'E'). The slices are then chosen sequentially from each of these five groups. After each trigger the lower slices of the brain are scanned first (during diastole) and the upper ones last. In this example one slice from each of the three groups labeled from (A) to (C) would be collected during diastole while two slices from groups (D) and (E) would be imaged, respectively, 570 and 760 ms after the trigger (during systole)—slices labeled 1–5 in Fig. 1. To avoid cross-talk effects, the slices are ordered in an interleaved manner within each of the groups. In the example chosen, the slices which would be imaged in the following cycle are labeled from 6 to 10 in Fig. 1. For both the non-gated and standard gated method, the slices were simply interleaved.

To correct for motion all data sets were re-aligned to a non-DW image using an affine registration method [11]. All diffusion-weighted images were then transformed into standard space after registration to each individual's structural scan in an intermediate step.

The standard deviation of the signal in each pixel was calculated over the 40 volumes to check for the presence of artifacts. An F test was performed to compare the variability observed with each of the three methods and a mask corresponding to the area covered for all data sets included in the analysis calculated. The threshold for the F-maps was set to $F = 5.6$ corresponding to a Bonferroni corrected p value of 0.05 (number of voxels in the common mask ~219,000).

To illustrate how the presence of pulsatile motion artifacts may influence the measurement of tensor-derived parameters, an additional experiment was performed. To be able to estimate the tensor, besides a non-diffusion-weighted image, images sensitized to diffusion along the following six directions were acquired: (1, 0, −1); (1, 0, 1); (0, 1, 1); (0, 1, −1); (1, 1, 0); (−1, 1, 0). For each type of acquisition (standard gated method, optimized gated method, and non-gated approach), 10 repeats of the full tensor data were collected, corresponding to a total of 70 volumes for each method. All the remaining imaging parameters were kept the same. This protocol was performed on four subjects (two males and two females). As before, to correct for motion all data sets were re-aligned to a non-DW image and the diffusion tensor was fit to each of the 10 sub-sets of seven images [12]. The fit was repeated after averaging together the 10 repeats acquired with each approach.

3. Results

Fig. 2 shows an example of an image which displayed pulsatile artifacts, together with the corresponding image acquired with the optimized gated scheme. Signal loss due to pulsatile brain motion is visible in the medial regions of the brain as indicated by the white arrow. Artifacts such as this were observed consistently for all the subjects on the data sets collected with no cardiac gating. A simple visual inspection was performed to evaluate what fraction of the images was affected. When considering all the repeats corresponding to a slice located in the affected area it was found that approximately 18% of the volumes displayed severe artifacts in the same consistent locations. A similar value of 20% had previously been reported in [6] for non-gated acquisitions. These values are consistent with the relative duration of systole compared to the whole cardiac cycle. Such results are not surprising considering that non-gated images are not acquired in any particular order in relation to the cardiac cycle. In contrast, no obvious artifacts due to pulsatile

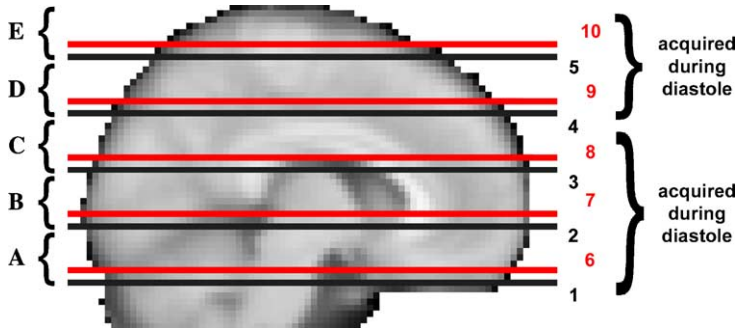

Fig. 1. Slice ordering using the optimized gating scheme. For a maximum estimated heart rate of 60 bpm, for example, five slices can be scanned per cardiac cycle. On our system, the effects of systole in the brain occur approximately 600 ms after the peripheral gating trigger. To ensure that the upper regions of the brain are imaged during systole, slices from each of the five sets (A–E) depicted are sequentially selected from the bottom (A) to the top of the head (E). Within each group the slices are ordered in an interleaved manner. To exemplify, the slices acquired during the first cardiac cycle are labeled in black, while those collected after the second trigger are shown in red. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this paper.)

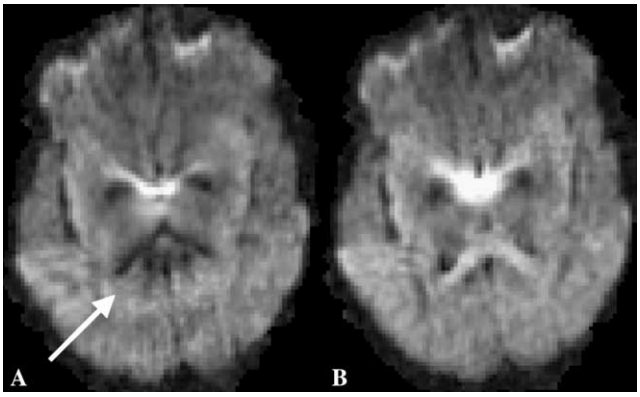

Fig. 2. (A) Image acquired without gating displaying artifact due to pulsatile motion. (B) Same slice imaged with the optimized gating scheme.

motion were visible in the data sets collected with the optimized scheme.

For the standard scheme, artifacts were detected on one of the data sets. This was probably due to the higher heart rate presented by this subject (84 bpm on average). For a heart rate of this range it would probably have been advisable to reduce the number of slices acquired per cycle while using the standard method. The acquisition of three slices requires 570 ms on our system, which corresponds to a significant portion of the cardiac cycle in this case (714 ms). The presence of artifacts implies that systole was not successfully avoided for this subject. However, the fact that no artifacts were detected in the data set acquired using the optimized scheme suggests that a more efficient use of time is possible. Artifacts can be avoided provided that the slices are acquired in the appropriate order. To avoid biasing the results towards the optimized approach, this subject was excluded from the rest of the analysis.

The presence of artifacts was detected by calculating the variance of the signal in each voxel over the 40 diffusion-weighed volumes. F tests were performed for the six remaining subjects to compare both the non-gated and the optimized gating schemes with the more standard gating method [6]. After thresholded for a Bonferroni corrected p value of 0.05, the F-maps were binarized and the results obtained for the six subjects summed together. A voxel value of 6 thus corresponds to significant differences being observed in all six subjects, whereas a value of 1 indicates that a significant difference was observed only in one case. The maximum intensity projections are shown in Fig. 3.

As can be seen in Fig. 3, voxels with increased variance were detected in the non-gated data sets, particularly in the case of regions of the brain below the corpus callosum, when compared to the standard gated acquisition. Differences in variance were consistently detected in areas such as the splenium of the corpus callosum, the brainstem, and the cerebellum. The genu of the corpus callosum was also affected, although the overlap between subjects was not as consistent.

When comparing the optimized gated sequence to the standard gated approach, it can be noted that although significant differences could also be detected, these did not occur systematically in the same voxels across subjects, but only ever in one subject at most. These voxels were located mainly in the upper areas of the brain. Looking at the equivalent area in the F-maps comparing the non-gated acquisition with the gated approach, a similar pattern can be visible in this part of the brain. Examining the raw images, no signal attenuation is visible in these areas, as would be expected if such differences were due to pulsatile brain motion. We therefore conclude that these differences can be explained by various experimental effects. For example, in some images signal arising from fat surrounding the brain could be detected, indicating incomplete fat-saturation. Note that although the fat is actually located around the brain, as water and fat possess different resonance frequencies,

the fat signal is spatially shifted on the images. Chemical shift artifacts could be seen in images acquired with any of the three methods, indicating that its presence is not related to whether gating was performed or not.

The better performance of the gated approaches relative to that of the non-gated method may be confirmed from Table 1. In this table the spatial mean of the standard deviation, calculated in the potentially affected areas, is shown. For this purpose a mask was used that contained the union of all voxels, when considering the six subjects, in which the variance in the non-gated case was significantly higher than when using either of the gated methods. The average acquisition times for each method are shown in Table 1. Note that the acquisition time for the non-gated approach was not always the same. This was because the number of slices acquired per subject could also be increased above the default value of 42, so as to be a multiple of the number of slices acquired per cycle when using the optimized method. The scanning efficiency of each method, i.e., the fraction of the total time that was actually spent acquiring data, is also shown in Table 1. In each case the acquisition time for the non-gated method was taken as reference. The scanning efficiency using the optimized gating method was significantly improved when compared with the standard gating method [6] ($p < 0.001$), allowing for a significantly higher SNR to be achieved per time unit. The average SNR for a single volume was determined by drawing similar regions of interest for the six subjects within a region of cortical gray matter. The mean values obtained are shown in Table 1 (mean ± standard deviation). No significant differences were observed between methods in this case.

In the DTI data sets acquired, a simple visual inspection resulted in the detection of artifacts in between 10 and 20% of the images acquired with no cardiac gating. No obvious artifacts were detected in the data sets acquired with cardiac gating.

When examining the non-gated data, it was possible to verify that the effect of pulsatile motion was more striking in the lower regions of the brain, such as the cerebellum or the brainstem although artifacts could also be identified in other areas of the brain such as the splenium of the corpus callosum.

This data also demonstrated that the effect of pulsatile motion is dependent on the orientation of the diffusion gradients relative to the direction of motion. For instance, when examining the cerebellum, artifacts could more easily be identified when the diffusion gradients had a significant component along the z axes, which again is in agreement with a superior–inferior direction for motion in this part of the brain.

To illustrate the potential effects of not gating on DTI quantities, one of the examples found in the DTI data acquired is shown in Fig. 4. This example corresponds to a situation where clear artifacts were detected in the cerebellum as shown in Fig. 4B. The color coded maps obtained for all three acquisition methods are shown in Fig. 4C. In the map obtained for the ungated acquisition, a higher anisotropy is visible in the region where the artifact occurred, while the direction of the principal eigenvectors appears to be biased towards the direction along which diffusion was being measured: $(1, 0, -1)$. The maps corresponding to both gated methods are, on the other hand, very similar. When the images corresponding to all 10 repeats were averaged together, it became difficult to detect any signal attenuation due to cardiac pulsation (bottom row of Fig. 4C). Under these circumstances the maps obtained for all three methods appear to be much more similar, although a small increase in anisotropy may still be visible in the ungated case.

4. Discussion and conclusions

As demonstrated in Fig. 3, consistent differences were observed when comparing the signal variance for the non-gated versus the standard gating method, particularly in regions located below the corpus callosum. In the study performed by Skare and Andersson using a 1.5 T scanner [6], only a few of the voxels located above the corpus callosum displayed significant differences. In our case, the number of voxels in this area where differences could be observed was increased, either when the non-gated or the optimized method were being compared to the standard approach. The exact location of these voxels was, however, not consistent among subjects. Inspection of the raw data revealed signal

Table 1
Comparison of the performance of the three acquisition methods used

	Non-gated acquisition	Optimized scheme	Standard method
Mean standard deviation	15 ± 5	9 ± 2	10 ± 4
Mean acquisition times	5 min 36 s ± 10 s	7 min 09 s ± 36 s	10 min 37 s ± 40 s
Scanning efficiency (%)	100	79 ± 8	54 ± 10
Average SNR	18 ± 3	18 ± 3	17 ± 2

Mean standard deviation across subjects in the affected areas (arbitrary units). Also shown are the mean acquisition times and scanning efficiencies of both gating schemes compared to the non-gated approach. The last row corresponds to the SNR measured for each method in a peripheral gray matter region. The average value observed for the 40 volumes acquired was determined and the mean across subjects calculated. The error bars correspond to the standard deviation across subjects.

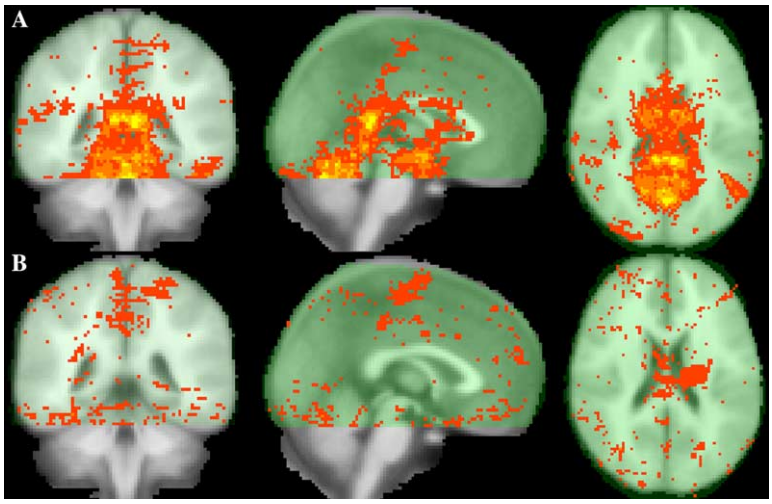

Fig. 3. Artifact reduction using cardiac gating. The voxels where the variance was significantly greater in the non-gated data set compared to the standard gating method are shown in (A), while the voxels where the optimized scheme led to a significantly higher variance than the standard gating approach are shown in (B). The thresholded, binarized *F*-maps obtained for all subjects were added together. A voxel value of six (bright yellow) therefore corresponds to significant differences being detected for all subjects, whereas a value of one implies that significant differences were detected only for one subject. The brain volume which was covered in all scans is depicted in green.

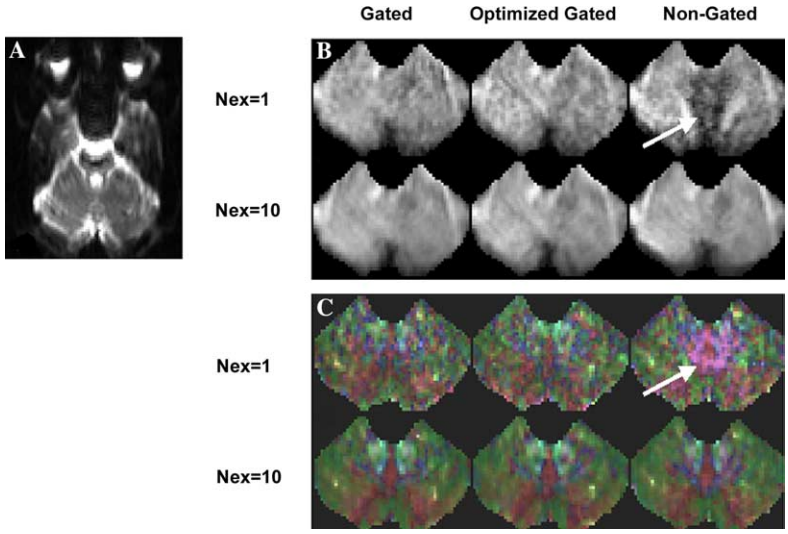

Fig. 4. (A) Non-diffusion-weighted image of a slice passing through the cerebellum. (B) The cerebellum was manually extracted from diffusion-weighted images of the same slice obtained for the diffusion orientation (1, 0, −1). On the top row individual images obtained with each acquisition method are shown. Increased signal attenuation can be seen in the ungated image as a result of cardiac pulsation (indicated by the white arrow). Averaging together the images obtained for the 10 repeats acquired, the presence of this artifact is masked out and its detection is no longer obvious. (C) Color-coded maps estimated from the DTI data sets corresponding to the images shown in (B). The intensity on these maps is determined by the estimated anisotropy whereas the color reflects the direction of the principal eigenvectors (red, left–right; blue, inferior–superior; green, anterior–posterior). The presence of the artifact in the ungated data led to an increased level of anisotropy and to bias in the orientation of the principal eigenvectors. As the diffusion-weighted image affected corresponded to the (1, 0, −1) gradient direction, purple is the dominant color in the region corresponding to the artifact seen in the ungated image (indicated by a white arrow). By averaging the data, the effects of the artifact were diluted and no obvious differences can be seen in the color-coded map obtained for the ungated data when compared to the maps for the gated acquisitions.

fluctuations in this area resulting from chemical shift artifacts. The greater impact of this type of artifact in our data can be understood by taking into account that the experiment was performed at 3.0 T. As it is harder to achieve a good B1-homogeneity throughout the whole brain at higher fields, the fat-saturation pulse is not as effective. For our scanner, such in-homogeneities affect mostly the top and bottom of the head where the effective RF pulses deviate more significant from their nominal value. Diffusion-weighted images are particularly susceptible to the chemical shift artifact as the diffusion coefficient of fat is much lower to that of brain tissue

which is why the signal arising from the fat is much less attenuated than that of the brain.

Another aspect which may explain why differences were observed between the two studies is that different approaches were taken to avoid excessively high thresholds when performing the Bonferroni correction. In the study presented in [6], the original data was resampled to a lower resolution, which resulted in an improvement of the SNR of the data and also reduced the number of voxels that needed to be considered when determining the threshold level F. The application of smoothing may, on the other hand, contribute to mask out some less prominent artifacts reducing the sensitivity of the F test. We therefore preferred to use a higher number of repeats (40 vs 18 used in [6]).

It is also possible that the use of a higher field scanner may contribute to increase sensitivity to brain pulsation. As discussed by Wirestam et al. [5], the amount of signal dephasing due to tissue translation should be dependent on the amplitude and duration of the gradient fields applied. The presence of susceptibility differences between, for instance, tissue and air located in the sinuses also results in the production of local gradient fields, the intensity of which is dependent on the strength of the static field. Throughout the cardiac cycle, nonlinear brain motion occurs such that the relative distances between different brain structures varies along the cycle. It is therefore possible, particularly in lower regions of the brain, where both susceptibility differences and pulsatile brain motion are most significant, that such susceptibility-induced gradients may contribute to additional dephasing and therefore to increased signal loss due to cardiac pulsation.

Several studies have demonstrated that the quality of the data are improved by performing cardiac gating [2,5–8]. However, the long acquisition times associated with traditional methods of cardiac gating still deter many researchers from applying them. Another aspect has to do with the effect of averaging the data. The occurrence of pulsatile motion artifacts on non-gated data sets depends on several factors. First, the artifacts only occur if an area of the brain which displays pulsatile motion happens to be imaged during systole. The frequency at which these artifacts appear should therefore increase with the heart rate of the subjects. The extent of their effects should also be variable, depending in particular on how much do the direction of motion and that of the diffusion gradients applied coincide. The effect of such artifacts on the estimates of anisotropy should also be dependent on the relationship between the preferred direction of motion and that of the local fibers. For instance, in an area of low anisotropy, the signal measured along different directions should normally be very similar. If, for one of the measurements, signal attenuation due to pulsatile motion is observed, then the anisotropy will be over-estimated. This effect is illustrated in Fig. 4. However, in the case of an area of high anisotropy, the effect may not always be the same. The calculated anisotropy may be overestimated or underestimated in the presence of the artifact, depending on whether the artifact occurs for a gradient direction that lies along the fiber direction or along another direction, respectively. Additionally, in the case of the main experiment described above, it should be noted that due to the already low SNR of voxels whose fibers lie along the z-direction, it may be difficult to accurately assess the benefits of cardiac gating in that subset of voxels.

The random nature of this artifact and the fact that its effects may also be variable imply that, when multiple repeats are simultaneously considered, the sporadic changes in FA and MD are diluted and may no longer be noticeable when average values are compared (see Fig. 4). The difficulty to detect these effects when averaging is performed may explain why many researchers still opt not to perform cardiac gating. It should be noted, however, that even if the effects of cardiac pulsation can be expected to be more subtle in this case, this does not mean that they are inexistent. This is particularly important if the data are to be used for fiber tracking as erroneous estimates of the principal direction of diffusion, even if only affecting a few voxels, may lead to significant differences in the reconstructed fiber tracts. On the other hand, as illustrated in Fig. 4, the effect on individual acquisitions may be quite significant, particularly if only a few diffusion directions are sampled. This is likely to be the case in a clinical setting, where limited scanning time may render impossible the acquisition of either a higher number of directions or averages.

A decisive factor which explains why cardiac gating still tends not to be used, even in this case, is the severe reduction in scanning efficiency when performing gating in the standard way, as no data are acquired during a significant fraction of time. If the acquisition times for the two approaches (gated and no-gated) were set to be equal, a higher number of samples would be obtainable in the non-gated approach. Since the SNR is proportional to the square root of the number of acquisitions, this would lead to an improvement in the SNR, provided that the quality of the two sets of images were comparable. However, Skare and Andersson [6] have pointed out that the potential reduction in noise achieved by doubling the acquisition time would still be lower than the increase in variability observed due to the presence of artifacts. Following this reasoning, the SNR achieved for an equivalent acquisition time without gating would actually be lower than when gating [6]. On the other hand, this argument should only be valid in areas of the brain which are deformed during the cardiac cycle while the SNR in non-affected areas should still increase. Not performing cardiac gating will therefore result in a spatially variable SNR. This disparity is undesirable as it may contribute to having

higher levels of confidence associated with fiber tracts reconstructed in the upper and more peripheral regions of the brain (where a higher SNR would be present) compared to those in regions affected by this kind of artifact.

To investigate the possibility of minimizing the level of artifacts in the images without excessively compromising scanning efficiency, an optimized method to perform cardiac gating was evaluated. It is important to note that in some cases modifications to this optimized method may be required. For instance, if the number of slices is increased so as to include the cerebellum, dividing the range of slices evenly according to the number of slices acquired per cardiac cycle may not be adequate. In this case, the number of slices belonging to each fraction may need to be revised, to ensure that none of the slices collected during systole are located below or at the level of the corpus callosum. The relative size of each fraction may therefore be different and, although a higher number of cardiac cycles will be required to complete the acquisition, the total acquisition time should still be improved compared with the standard gating method. A similar approach may need to be followed when scanning subjects with high heart rates. As the number of slices acquired per cycle diminishes, the scanning range will be divided into a smaller number of fractions which will each cover a larger extent of the brain. Since the upper fractions may start to include artifact-prone areas, a re-distribution of the slices will be necessary. Furthermore, in the particular MR system used in this work the interval between the peripheral trigger and the observation of the highest level of variance in the data (600 ms) led to selecting the slices from each of the groups described starting from the bottom to the top of the brain. However, a different time delay of 100 ms had also previously been reported [6], which indicates that it is likely that the introduction of extra hardware delays may lead to differences in the optimal slice ordering from system to system. This is because the ordering should be done so as to ensure that the upper slices of the brain are the ones to be acquired during systolic motion, regardless of the actual timings. The implementation of the optimized method may be straightforward in some systems, provided that the most favorable situation corresponds to selecting a slice from each block sequentially either from the bottom to the top of the brain or vice versa, as this option is more likely to be available.

The optimized method should work best for subjects with steady heart rates. To maximize scanning efficiency, the number of slices collected per cycle is calculated so as to reduce the waiting period between triggers. If the duration of the cycle becomes shorter than the time needed to acquire the selected number of slices, the next trigger is missed and the efficiency reduced. To avoid this situation, a safety margin can be added to the estimate for the maximum heart rate for that subject. This is particularly important when the subject's heart rate is close to a transition value, above which a lower number of slices can be scanned. In this case the estimate can be increased so that it exceeds the closest transition value. On the other hand, if the heart rate is significantly overestimated, large gaps will exist between acquisition periods and the scanning efficiency will not be as high as possible.

Another important issue when performing cardiac gating is the possibility of signal changes due to variable levels of T_1 weighting. As the subject's heart rate varies, so does the effective repetition time TR, leading to fluctuations in the amount of T_1-weighting. However, as a long TR can be used for whole brain acquisitions without loss of efficiency, such fluctuations become insignificant. In the approach described the effective TR was always at least 8.0 s, which is significantly higher than the average T_1 in human brain white matter at 3.0 T (860 ms) [13]. To investigate the possible effects of heart rate fluctuations, simulations were performed considering baseline heart rates from 40 to 120 bpm. The number of slices acquired per cycle and the expected repetition time TR were determined for each of the heart rates considered (increments of 1 bpm). The effective repetition times corresponding to fluctuations of -10 to $+10$ bpm over each baseline heart rate were calculated. Considering the signal to be proportional to $[1 - \exp(-\mathrm{TR}_{\mathrm{eff}}/T_1)]$, the relative signal changes $(S_{\mathrm{eff}} - S_{\mathrm{bas}})/S_{\mathrm{bas}}$, where S_{bas} corresponds to the baseline signal (and repetition time TR), were calculated for each situation. Following these procedures, the maximum absolute signal change for white matter was found to be 0.009%. This change represents a very small fraction of the baseline signal, much lower than the fluctuations expected to occur simply due to the presence of noise and can therefore safely be ignored.

In summary, we have demonstrated that it is feasible to acquire cardiac-gated DW images in a more efficient way, provided that the slices are ordered in a particular manner. Using this approach, images free from cardiac pulsation artifacts can be acquired at a very small cost in acquisition time, when compared with more artifact prone non-gated acquisitions.

Acknowledgments

This work was supported by the Calouste Gulbenkian Foundation (R.G.N.), the UK Medical Research Council (S.C.), and the Dunhill Medical Trust (P.J.).

References

[1] A.W. Anderson, J.C. Gore, Analysis and correction of motion artifacts in diffusion weighted imaging, Magn. Reson. Med. 32 (1994) 379–387.

[2] C. Pierpaoli, S. Marenco, G.K. Rohde, D. Jones, A. Barnett, Analyzing the contribution of cardiac pulsation to the variability of quantities derived from the diffusion tensor. in: Proceedings of the 11th Annual Meeting of the ISMRM, Toronto, 2003. p. 70.
[3] O. Dietrich, S. Heiland, T. Benner, K. Sartor, Reducing motion artefacts in diffusion-weighted MRI of the brain: efficacy of navigator echo correction and pulse triggering, Neuroradiology 42 (2000) 85–91.
[4] D. Greitz, R. Wirestam, A. Franck, B. Nordell, C. Thomsen, F. Stahlberg, Pulsatile brain movement and associated hydrodynamics studied by magnetic resonance phase imaging. The Monro-Kellie doctrine revisited, Neuroradiology 34 (1992) 370–380.
[5] R. Wirestam, D. Greitz, C. Thomsen, S. Brockstedt, M.B. Olsson, F. Stahlberg, Theoretical and experimental evaluation of phase-dispersion effects caused by brain motion in diffusion and perfusion MR imaging, J. Magn. Reson. Imaging 6 (1996) 348–355.
[6] S. Skare, J.L. Andersson, On the effects of gating in diffusion imaging of the brain using single shot EPI, Magn. Reson. Imaging 19 (2001) 1125–1128.
[7] R. Turner, D. Le Bihan, J. Maier, R. Vavrek, L.K. Hedges, J. Pekar, Echo-planar imaging of intravoxel incoherent motion, Radiology 177 (1990) 407–414.
[8] D.K. Jones, C. Pierpaoli, Contribution of Cardiac Pulsation to Variability of Tractography Results. in: Proceedings of the 13th Annual Meeting of the ISMRM, Miami Beach, 2005, p. 222.
[9] S. Brockstedt, M. Borg, B. Geijer, R. Wirestam, C. Thomsen, S. Holtas, F. Stahlberg, Triggering in quantitative diffusion imaging with single-shot EPI, Acta Radiologica 40 (1999) 263–269.
[10] T.G. Reese, O. Heid, R.M. Weisskoff, V.J. Wedeen, Reduction of eddy-current-induced distortion in diffusion MRI using a twice-refocused spin echo, Magn. Reson. Med. 49 (2003) 177–182.
[11] M. Jenkinson, P. Bannister, M. Brady, S. Smith, Improved optimization for the robust and accurate linear registration and motion correction of brain images, Neuroimage 17 (2002) 825–841.
[12] P.J. Basser, J. Mattiello, D. LeBihan, Estimation of the effective self-diffusion tensor from the NMR spin echo, J. Magn. Reson. B 103 (1994) 247–254.
[13] S. Clare, P. Jezzard, Rapid T(1) mapping using multislice echo planar imaging, Magn. Reson. Med. 45 (2001) 630–634.