Introduction

Anyone who has performed an MR diffusion experiment has discovered that the images look quite different from structural MR images, the latter appearing almost like photographs taken after someone has cryosectioned the object.

The difference in contrast aside the diffusion images look strangely distorted with some parts stretched out like rubber sheets and other parts squashed together creating ``thick piles'' of MR signal.

On closer inspection it becomes apparent that the distortions are not evenly distributed, but rather appear mainly in areas adjacent to air filled cavities. Furthermore the distortions run only along the phase-encode direction.

The purpose of this talk is to explain why you are seeing this, and also to give advice about how to avoid and correct for these problems.

What is an off-resonance field?

An ``off-resonance field'' is simply an "image" that depicts the difference between what we think the field is, and what it really is. In a 3T scanner we would expect the field to be 3T at any point well inside the bore, but in reality it might be closer to 2.9999T making the off-resonance field 0.0001T. The reason it is called a ``field'' is because this difference will not be the same at every point inside the scanner, therefore requiring an image where each voxel value signifies the difference between the expected and actual field at that point.

Typically the field is defined in units of Hz, which means that the value is the difference in Larmor frequency for a water molecule in the expected and the actual field. Typical values are in the order of tens of Hz, which means that the deviations are in the range of ppm of the main magnetic field.

Where does the off-resonance field come from?

There are several sources of off-resonance field, but this talk only addresses that caused by "susceptibility". If you look up the word susceptible you will find something like "likely or liable to be influenced or harmed by a particular thing" and in our context that "thing" is the magnetic field.

Different materials are differently susceptible to become magnetized and for our purposes the main "materials" are tissue and air. Air is more susceptible to become magnetized than tissue so in an air filled cavity the field will be slightly stronger than in tissue. For brain imaging such cavities would be for example the ear canals and the sinuses.

The resulting field will *not* just be a scaled version of the susceptibility map as it (the field) must satisfy Maxwell's equations. That means that even if we had exact knowledge of the external field and the susceptibility map it would be non-trivial to calculate the resulting field.

Why are EPI images so sensitive to the off-resonance field?

Every time we stick an object in the scanner there will be a non-trivial off-resonance field and yet most MR images are anatomically faithful, so what is the problem with EPI?

All MR imaging is based on coding location by frequency, i.e. the frequency of the signal we detect determines the location it came from. The difference in frequency between adjacent pixels is known as the bandwidth (BW) per pixel. The smaller the BW per pixel the more sensitive the sequence is to off-resonance.

If the BW is 10Hz an off-resonance value of 10Hz results in a one pixel displacement (distortion) of the signal. If on the other hand the BW was 1000Hz the same off-resonance would yield a negligible (0.01pixel) displacement.
An EPI image has a vastly lower BW (~10Hz) in the (unfortunately named) phase-encode direction compared to in the frequency-encode direction (and also compared to most other sequences). This results in an image that is severely distorted along one direction but not along the other.

Can we make our EPI sequence less sensitive?
The main advantage of parallel (in-plane acceleration, IPAT) EPI imaging is that it increases the BW per pixel. If you for example use GRAPPA with an acceleration of three your BW per pixel will increase by the same factor and your distortions will be reduced by a factor of three. However, IPAT also have its cons, as it will reduce the SNR. Therefore IPAT should be combined with retrospective correction techniques so as to obtain the optimal trade-off between loss of SNR and reducing distortions to a level where they are feasible to correct.

Can we correct the images if we know the off-resonance field?
If we know the off-resonance field we also know to where the signal has been displaced. When we want to know the "true" signal for a given voxel we can look at the off-resonance field for that voxel and that tells us how far we need to move in the acquired (distorted) image to find that "true" signal. Hence, we can *in principle* correct the distorted image if we know the field.
However, in areas where the field is changing rapidly the signal "rightfully belonging to" several voxels can get squashed into fewer, or even a single, voxels. Therefore there is a limit to how well the images can be "fixed" retrospectively, which is the reason one might want to combine it with some IPAT technique.

How can we know the off-resonance field?
Techniques for finding the field can be broadly divided into those that attempt to measure it directly and those that calculate it from two or more images that are differently distorted.
The quintessential example in the former category is a "fieldmap sequence", which typically consists of two gradient-echo images acquired with slightly different echo times. If all voxels experience exactly the resonance frequency the two (complex) images will have the same phase. If, on the other hand, the true frequency in a given voxel is slightly different that frequency will lead to the accumulation of a phase that is different in the two images. The off-resonance frequency can be calculated by dividing that phase-difference by the echo time difference.
The other category is typified by a pair of images with opposing phase-encode directions, which leads to distortions of equal magnitude but in opposing directions. The field can be calculated from such a pair by iteratively trying to find the field that when applied to (used to correct) both images in the pair produces two corrected images that are as similar as possible.

Correcting the images – revisited.
An advantage of the latter technique is that for any area where the signal has been squashed together in one of the images (leading to a loss of information) it

has been stretched in the other image. That means that there is in principle no irretrievable loss of information and it should be possible to achieve a better retrospective correction.

Combining corrections.

The susceptibility-induced distortion is not the only source of spatial artifacts that affect diffusion imaging. We also need to handle eddy current-induced distortions and subject movement. Even if we, as a first approximation, assume that the susceptibility-induced field remain constant w.r.t. subject movement these source interact with each other. The susceptibility-induced field remains constant in the object frame of reference (though the direction of the distortion is not) whereas the eddy current-induced field is fixed in the scanner frame of reference.

This begets the question of what order to apply the corrections in? I believe they should be applied together in a framework that takes those interactions into consideration. The last part of the talk will briefly introduce such a framework.