More Than Motion – Simultaneous Motion, Eddy Current & Susceptibility Correction

Speaker Name: Jesper Andersson
Email: jesper.andersson@ndcn.ox.ac.uk
Organization/employer: University of Oxford
Address:
FMRIB Centre
University of Oxford
John Radcliffe Hospital
Oxford OX3 9DU
UK

This talk will describe a method to correct for eddy current-induced distortions and, of particular interest for this audience, subject movement in diffusion weighted images.

Diffusion imaging is typically performed by acquiring multiple spin-echo EPI images preceded by diffusion sensitising gradients that spoil the signal in proportion to the mobility of spins (i.e. water). Current trends in research oriented diffusion studies is towards large numbers of diffusion directions and/or multiple *b*-values. It is not unusual to acquire a few hundreds of volumes for a study.
Each volume is preceded by a unique diffusion gradient (defined by direction and magnitude) which will elicit eddy currents and those will cause an off-resonance field during the image encoding. This will result in distorted image volumes, with each volume distorted in a unique way, due to the sensitivity of EPI to off-resonance. In addition to that the relatively long acquisition times makes it almost inevitable that there will be subject movement.
The off-resonance fields are spatially smooth and slowly varying which means that they can be described by a small number of parameters. Hence it would superficially seem like a trivial problem to register the images together using "traditional" methods. However, it is complicated by the fact that in addition to being differentially distorted and being in different positions the images have inherently different contrast. Especially for high *b*-values two volumes acquired with different diffusion directions can be virtually unrecognisable as being from the same subject.

We have attempted to solve this by making predictions, for each unique volume, what we "think the undistorted volume should look like" and then aligning the acquired volume to the corresponding prediction. For this to work hinges on the "average" space of the acquisitions being distortion free. A simplistic example of how such a prediction could be made would be to have acquired volumes with the diffusion gradient **g** and the gradient **-g**, which would result in two images with the same diffusion contrast but with different (opposing if ones assumes linearity) distortions. If we now "predicted" that the "true" image would be the average of those two images we would have a prediction in (approximately) undistorted space, albeit a little blurred as the two images are in different spaces. One could then register both images to this "blurry prediction", thus nudging them both towards undistorted space. The next prediction would be the average of the two "nudged" images, yielding a less blurred image in undistorted space. After a few iterations of the register-predict cycle we would expect the images to be registered.

The simplistic "prediction maker" above suffers from a number of disadvantages, notably the need to acquire both **g** and **-g** (which is redundant in terms of diffusion information) and the poor SNR of predictions based on just two acquired volumes. Hence we instead use a Gaussian Process (GP) as a prediction maker. A GP can be thought of as a continuous extension of a multivariate Gaussian distribution (MVN). An n-dimensional MVN is characterised by an nx1 mean-vector and an nxn covariance matrix. If we have a partial sample, i.e. one where we have observed n-1 variables but are missing the n'th, we can often make a quite good prediction of the n'th variable. To make it concrete imagine a 2D MVN characterising height and weight in some population, and imagine that the mean vector is $[180\ 80]^T$ (cm and kg). If asked to guess the weight of some random person our best guess would be 80kg, and it would be associated with a relatively large uncertainty. If on the other hand the height was already known, to be e.g. 200cm, we could make a better guess, maybe 100kg, associated with a smaller uncertainty. A GP is similar in that it describes a stochastic process using a covariance, but since it is a continuous process it uses a covariance-function to specify the covariance between any two points from that process. So when “predicting” what the signal value should be for a specific diffusion gradient it uses the observed values for all gradients (in contrast to the two in our toy example above).

Given the above we can think of the Gaussian process as a prediction maker that given the whole data set can make a prediction about any given volume in a space that is closer to undistorted space than the corresponding acquired volume. The registration algorithm can thus be described as

```
set p = 0 // All distortion parameters zero
while (not converged) {
   for (every volume i) {
      undistort volume i given pi
      load undistorted volume i into prediction maker
   }
   for (every volume i) {
      get prediction for volume i
      distort prediction i given inverse of pi
      compare prediction to actual acquisition for volume i
      use Gauss-Newton to update pi given comparison
   }
   calculate if converged
}
```

These concepts will be described in greater detail and it will be shown that with this strategy one can successfully register images with very high b-values (we have tested up to b=10000s/mm^2).