Image quality transfer: exploiting bespoke high-quality data to enhance everyday acquisitions

Daniel C. Alexander[1], Darko Zikic[2], Viktor Wottschel[3], Jiaying Zhang[1], Hui Zhang[1], and Antonio Criminisi[2]
[1]*Dept. Computer Science, University College London, London, London, United Kingdom,* [2]*Microsoft Research, Cambridge, United Kingdom,* [3]*Institute of Neurology, University College London, London, United Kingdom*

Target Audience: Primarily diffusion MRI analysis researchers, although the method extends to other modalities.
Purpose: To exploit high-quality but hard-to-obtain images to enhance lower quality more readily-available data.
Introduction: The Human Connectome Project (HCP) [1] uses a bespoke scanner and ~1h acquisition time per subject to obtain in-vivo diffusion MRI data of unprecedented information content. We present a method to propagate information from such unique data sets to more everyday acquisitions and demonstrate in two ways: i) enhancing the spatial resolution of diffusion tensor images (DTIs); ii) enhancing the information content by predicting neurite orientation density and dispersion imaging (NODDI) [2] indices (normally requiring two b-values) from single-shell, b=1000s/mm^2, DTI data. The former is important for tractography, where image resolution is a key factor [1]; the latter, because NODDI normally requires two b-values [2] so is not possible on historical single-shell data sets.
Method: We learn a mapping from low-quality to high-quality image patches using HCP data. To enhance DTI resolution the input to the mapping is an $n\times n\times n$ cube (patch) of low-resolution voxels each containing a DT. The output is an $m\times m\times m$ high-resolution patch corresponding to the central voxel of the input patch; see figure 1. Matched pairs of training samples come from the b=1000s/mm^2 shell only: input vectors from downsampling raw images and fitting the DT; output vectors from fitting the DT at full resolution. To predict NODDI maps, the input is a patch of the full-resolution DTI (from b=1000s/mm^2 shell only); the output is the set of NODDI parameters (neurite density, f_{ICVF}; dispersion index, ODI; free-water fraction, f_{ISO}; and mean orientation) fitted to the full HCP data set (b=1000, 2000, and 3000s/mm^2). We use random-forest regression [3] to learn the mapping and compare to linear regression; see [4].

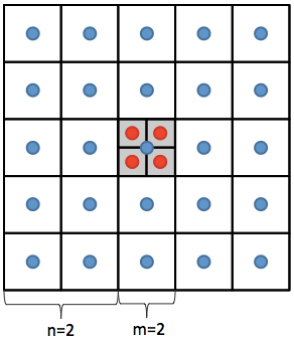

Fig. 1: Patch structure for super-resolution, which maps low-resolution patches of $(2n+1)^3$ voxels (blue dots) to high-res patches of m^3 voxels (red dots) in the central low-resolution voxel. The mapping is $6(2n+1)^3$ dimensions to $6m^3$, as DTs have six elements.

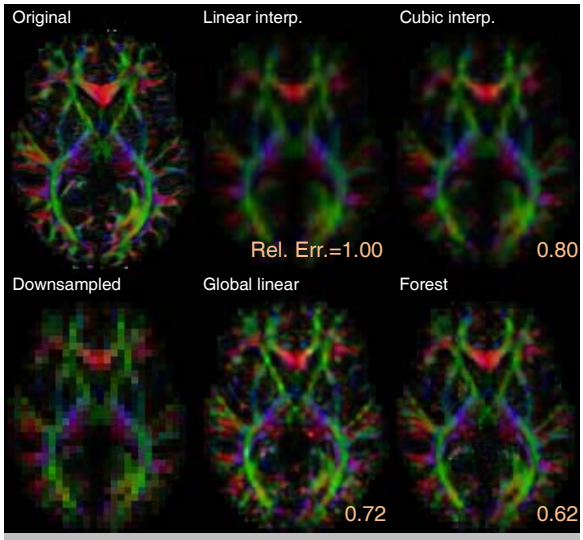

Fig. 2: Super-resolved direction-encoded colour FA maps from a low-resolution input image (bottom left) using standard interpolation (top right) and regression (bottom right); original full resolution image (top left). Here m=4 and n=2. Relative errors in orange (see text).

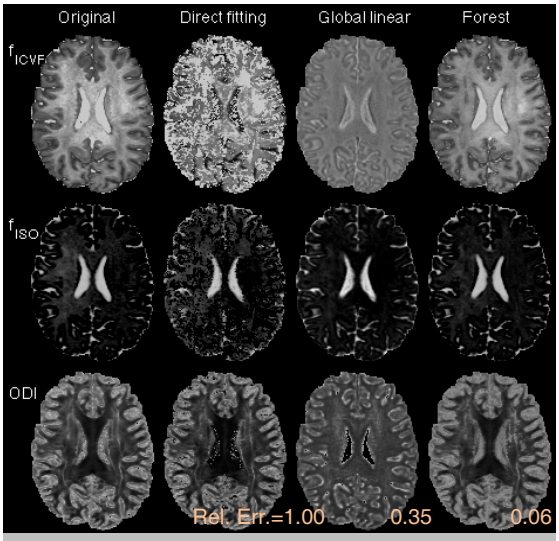

Fig. 3: NODDI parameter maps from direct model-fitting to the full raw data set (left) and b=1000 shell only (left middle), compared to those from linear (right middle) and random-forest (right) regression. Here m=1 and n=1. Errors average over the three indices.

Results: Figure 2 compares a ground-truth DTI (an HCP data set not in the training set) to reconstructions from a down-sampled image using each regression model and standard interpolation algorithms. Figure 3 uses the same unseen HCP data set to compare reconstructed NODDI parameter maps with ground truth. The relative error scores (orange) are root-mean-square voxel-to-voxel differences averaged over 8 unseen HCP data sets, as in [4], and normalised by the worst case ("Linear Interp" and "Direct Fitting") to quantify the improvement other techniques provide. Figure 4 demonstrates super-resolution and NODDI reconstruction from a low-resolution non-HCP data set acquired in 10 minutes with a 30-direction b=1000s/mm^2 clinical DTI protocol on a standard 3T scanner.
Discussion: Learned super-resolution outperforms standard interpolation, reducing errors to ~60%, potentially offering major benefits for tractography. The method is simpler than related techniques, e.g. [5], requiring only off-the-shelf computational tools; it is computationally efficient; and it requires no a-priori registration or segmentation. Recovery of NODDI parameters from b=1000s/mm^2 data alone fails using direct fitting [1] but the learned mapping predicts the parameters well. The non-linear "Forest" mapping improves on linear in both applications, although the learned mappings are compact with diminishing returns above 8 training data sets and 8 trees in the forest. Generalization ability is promising: performance on images from male subjects is almost identical with all male, all female, or mixed training sets. Further work will assess generalization among more diverse groups and in pathology. The method extends naturally to more complex voxel content, e.g. multi-fibre models etc, and to other image modalities.

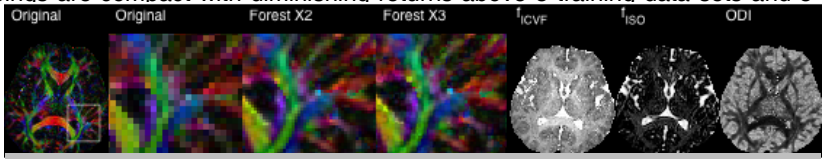

Fig. 4: Super-resolution and NODDI reconstructions from a non-HCP data set.

Refs: 1. Zhang NIMG 2012. 2. Sotiropoulos NIMG 2013. 3. Criminisi Springer 2013. 4. Alexander MICCAI 2014. 5. Coupé NIMG 2013