Provided for non-commercial research and educational use only.
Not for reproduction or distribution or commercial use.

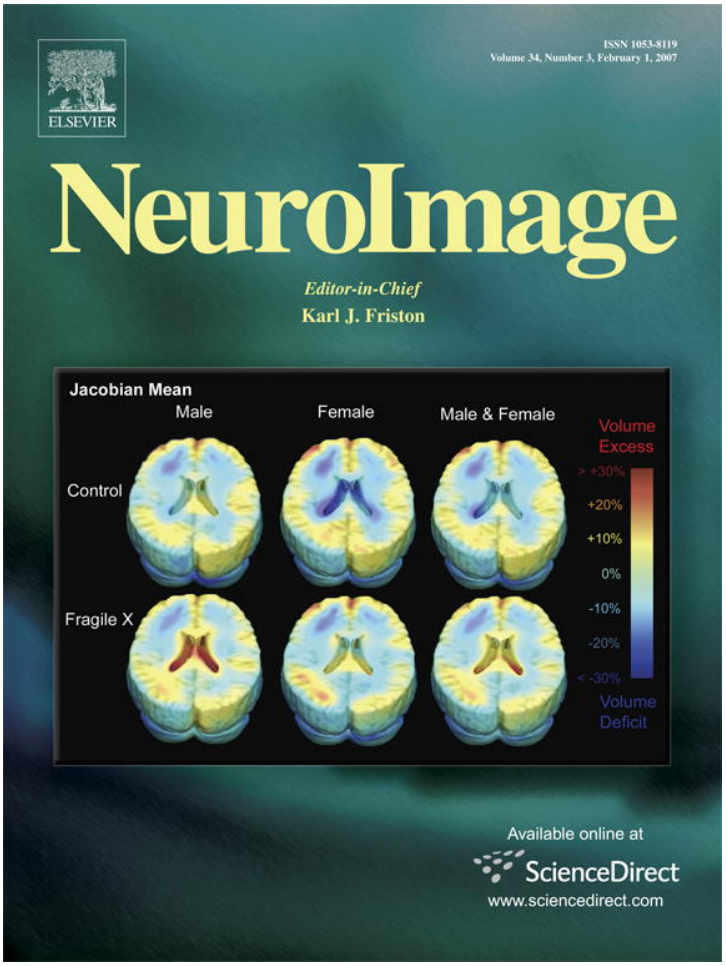

This article was originally published in a journal published by Elsevier, and the attached copy is provided by Elsevier for the author's benefit and for the benefit of the author's institution, for non-commercial research and educational use including without limitation use in instruction at your institution, sending it to specific colleagues that you know, and providing a copy to your institution's administrator.

All other uses, reproduction and distribution, including without limitation commercial reprints, selling or licensing copies or access, or posting on open internet sites, your personal or institution's website or repository, are prohibited. For exceptions, permission may be sought for such use through Elsevier's permissions site at:

http://www.elsevier.com/locate/permissionusematerial

ELSEVIER

NeuroImage

www.elsevier.com/locate/ynimg
NeuroImage 34 (2007) 1126 – 1135

The effect of large veins on spatial localization with GE BOLD at 3 T: Displacement, not blurring

Cheryl A. Olman,[a,*] Souheil Inati,[b] and David J. Heeger[c]

[a]Departments of Psychology and Radiology, University of Minnesota, N218 Elliott Hall, 75 East River Road, Minneapolis, MN 55455, USA
[b]Center for Neural Science and Department of Psychology, New York University, USA
[c]Department of Psychology and Center for Neural Science, New York University, USA

Received 3 May 2006; revised 30 August 2006; accepted 31 August 2006
Available online 6 December 2006

We used two different methods of region of interest (ROI) definition to investigate the spatial accuracy of functional magnetic resonance imaging (fMRI) at low and high spatial resolution. The "single-condition localizer" consisted of block alternation between a target stimulus and a mean gray background. The "differential localizer" consisted of block alternation between the target stimulus and another stimulus that filled the complement of the visual field. A separate series of scans, in which the target stimulus was presented briefly with long inter-stimulus intervals, was used to measure the hemodynamic impulse response function (HIRF). As expected, the differential localizer defined more restricted ROIs that better matched the predicted cortical representation of the target stimulus. However, at low resolution (3-mm isotropic) many voxels that responded positively to the target stimulus in the differential protocol responded negatively to the target stimulus in the single-condition localizer and in the HIRF measurements. The localization errors were attributed to voxels near large veins, which were identified based on low mean intensity and high variance. At high resolution (1.2-mm isotropic), the effects of large veins were present, but affected a smaller number of voxels. Thus, the use of differential localizers does not necessarily result in a more accurate indication of the underlying neural activity. Localization errors are reduced at higher spatial resolutions and can be eliminated by identification and removal of voxels dominated by large veins.
© 2006 Elsevier Inc. All rights reserved.

Introduction

Our ability to characterize human brain function is in many ways limited by the spatial resolution of blood oxygenation level-dependent (BOLD) fMRI. BOLD is the dominant technique for measuring and mapping neural activity in the human brain, but the effective spatial resolution of the technique is not well characterized. Even though the BOLD response is a spatially and functionally indirect measure of neural activity, many studies have demonstrated that BOLD experiments can provide excellent localization of neural activity (e.g., Cheng et al., 2001; for review see Heeger and Ress, 2002). To characterize the precision of the BOLD response, several studies have measured approximate cortical point-spread functions (Engel et al., 1997; Parkes et al., 2005). However, the accuracy with which BOLD fMRI can localize neural responses is known to be dependent on both experiment design and data analysis.

Functional MR image resolution has been improved significantly with advances in MRI technology. With the increased availability of high field systems, fast head-only gradient coils, and development of parallel imaging methods (Pruessmann, 2004), 1-mm image resolution is now readily achievable. The true spatial resolution of BOLD fMRI, however, is degraded by correlation between neighboring regions of cortex (Grinvald et al., 1994), $T_2^{(*)}$ blurring of the images (Farzaneh et al., 1990), and the presence of large draining veins (Duvernoy et al., 1981). While spin echo (SE) techniques can reduce the contribution of signal from large veins on the pial surface of the cortex (Yacoub et al., 2003), the gain in spatial sensitivity comes at the cost of a significant loss in signal-to-noise ratio (Norris et al., 2002; Parkes et al., 2005). While ultra high-field experiments can clearly benefit from the increased spatial specificity of SE BOLD (Yacoub et al., 2005), the low signal-to-noise ratio at 3 T means that the technique can fail to yield useful data. While conventional gradient echo (GE) BOLD consistently yields robust and reliable measurements, deoxyhemoglobin concentration at a given location in a moderately large vein can reflect changes in neural activity several millimeters away. Thus, where the BOLD signal is dominated by a large vein, the effective spatial resolution is poor regardless of the size of the volume element (voxel).

Many experiments for which uniformly good spatial specificity is important employ GE BOLD in combination with a differential experimental protocol (Cheng et al., 2001). Whereas the baseline in a single-condition protocol is resting state (no stimulus), the baseline in a differential protocol is the hemodynamic response evoked by a stimulus or stimuli designed to activate a complementary set of neurons (Bonhoeffer and Grinvald, 1993).

* Corresponding author. Fax: +1 612 626 2079.
E-mail address: caolman@umn.edu (C.A. Olman).
Available online on ScienceDirect (www.sciencedirect.com).

1053-8119/$ - see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2006.08.045

Author's personal copy

In the specific case of early visual cortex, complementary stimuli might occupy adjacent regions in the visual field, eliciting neural activity in adjacent regions of cortex. When adjacent regions of cortex are stimulated alternately, deoxyhemoglobin concentrations are not modulated in veins that receive blood from both cortical regions. Thus, at the boundary between such adjacent cortical regions, deoxyhemoglobin concentrations in large veins are not modulated and only small intracortical veins and capillaries contribute to BOLD contrast. Therefore, the accuracy with which BOLD fMRI can localize cortex responsive to the target stimulus should be better for a differential than for a single-condition protocol.

Another approach to improving the spatial accuracy of an fMRI experiment is to identify voxels in which the BOLD signal is dominated by large veins and exclude these from the analysis. In GE BOLD, voxels containing large veins have low intensity due to the short T_2^* decay time that results from a high local concentration of deoxyhemoglobin; fluctuations in blood flow and volume contribute to high variance in these same voxels (Duyn, 1995; Lee et al., 1995). Therefore mean-normalized variance can be used to identify which voxels are likely to have BOLD contrast dominated by large veins (de Zwart et al., 2005).

The experiments described here assess the ability of BOLD fMRI to localize visually evoked neural activity in primary visual cortex (V1). If a differential protocol minimizes the contribution of large draining veins, then the cortical activation map derived from a differential protocol should be a more accurate reflection of underlying neural activity than that derived from a single-condition protocol. If large draining veins are largely responsible for the spatial inaccuracy in the BOLD signal, then identifying and removing voxels dominated by large draining veins should improve localization accuracy. We tested these predictions by measuring the accuracy of BOLD fMRI with two imaging resolutions—3-mm isotropic voxels and 1.2-mm isotropic voxels.

Methods

Subjects

Four subjects (2 female, age 27 to 34) participated in the experiments, all of whom had normal or corrected to normal vision. The experimental protocols conformed to safety guidelines for MRI research and were approved by the Institutional Review Board at New York University.

Altogether, each subject participated in five scanning sessions: (1) retinotopic mapping, (2) high-resolution hemodynamic impulse response (HIRF) and localizers, (3) low-resolution HIRF and localizers, (4) high-resolution localizers, and (5) 3D MP-RAGE volume anatomy. The HIRF scanning sessions included two repetitions of each type of localizer, one of each type before and the other pair after the HIRF scans. The high-resolution localizer scanning session included 4 repetitions of each of the 2 localizers, differential and single-condition, at high spatial resolution.

Visual stimuli and experimental protocol

Stimuli were generated in Matlab (MathWorks, Inc., Natick, MA) and displayed by an EIKI LC-XG199 LCD projector with a custom zoom lens, housed outside the magnet room, onto a screen positioned behind the subjects' heads. Subjects viewed the screen through a mirror over their eyes. The target stimulus was a flickering checkerboard restricted to two annuli, extending from 1 to 3° and from 6 to 8° eccentricity (Fig. 1). The complement stimulus filled the complement of the visual field: a flickering checkerboard that was restricted to the center (0–1°), and two annuli (3–6 and 8–12°). Both stimuli were presented with a 75% duty cycle (750 ms out of every second) to minimize contrast adaptation, perceptual fading, and filling-in.

Three types of scans were run. In the "single-condition localizer", the target stimulus was presented for 8 s and then a mean gray screen was presented for 8 s, alternating for 10 ½ cycles or a total of 168 s. In the "differential localizer", the target stimulus was presented for 8 s and then the complement stimulus was presented for 8 s, again for 10 ½ cycles (168 s). In the "HIRF scans", the target stimulus was presented briefly (750 ms) every 20 s, 15 times per scan with 5 s before the first stimulus onset and 19.25 s after the last stimulus offset, for a total scan duration of 305 s. During every scan, subjects performed a task at fixation to control their attention and cognitive state. The white fixation point dimmed with random timing, on average every 3 s, and subjects were required to press a button when they detected the dimming. A 3-down, 1-up staircase governed the magnitude of the dimming, to

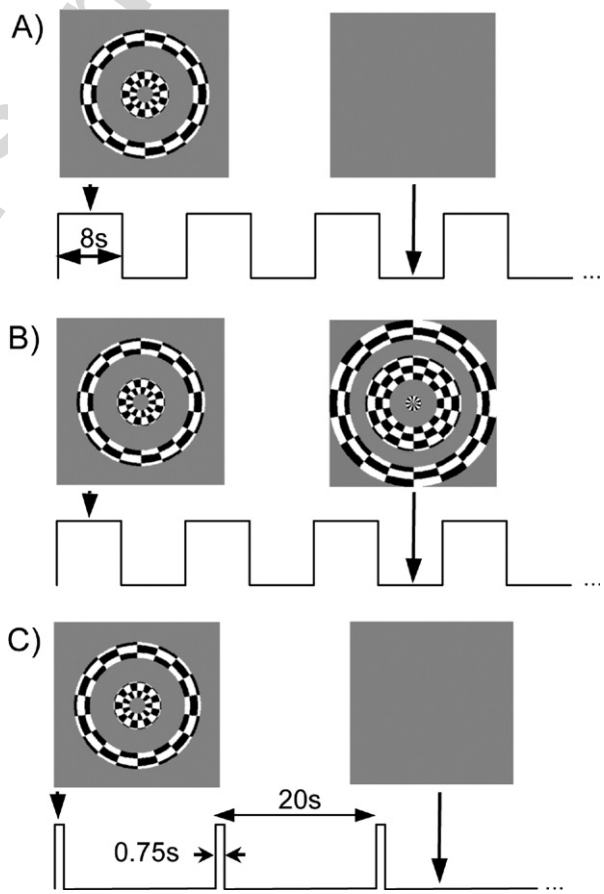

Fig. 1. Stimuli and stimulus presentation protocol. (A) Single-condition localizer: the target stimulus (a checkerboard flickering at 4 Hz, restricted to a pair of annuli subtending 1–3 and 6–8° of visual angle) was alternated with a mean-gray background in 8-s blocks. (B) Differential localizer: the target stimulus was alternated against its complement in the visual field in 8-s blocks. (C) Hemodynamic impulse response measurement: the target stimulus was presented for 750 ms, every 20 s.

Author's personal copy

maintain performance close to 80% and motivate constant effort and alertness.

Functional MRI

Data were acquired using echo-planar imaging (EPI) in an oblique axial orientation, parallel to and centered on the calcarine sulcus. Experiments were carried out on a 3 T Allegra scanner (Siemens, Erlangen, Germany) equipped with a volume transmit coil and a four-channel occipital surface receive array (NM-011 transmit head-coil and NMSC-021 receive-coil, Nova Medical, Wakefield, MA, USA). A 2D, navigated, single-shot, zoomed field of view EPI pulse sequence with interleaved slice ordering was used for all fMRI scans. This sequence is a 2D, single-shot version of that reported by Fleysher et al. (2005). To enable high-resolution acquisition with a TE of 30 ms, the field of view was reduced in the phase encode direction (anterior–posterior) and the signal from the front of the head was suppressed using outer volume suppression (OVS), in a manner similar to that reported by Pfeuffer et al. (2002). Each TR therefore consisted of a slab-specific saturation pulse for OVS, a frequency selective fat saturation pulse, a slice selective excitation pulse, three navigator echoes, a ramp-sampled rectilinear EPI trajectory, and spoiler gradients after the readout. The volume transmit coil used in these experiments provided a spatially uniform B_1 field, and allowed for efficient OVS. The raw data were reconstructed off-line using custom written C and Matlab code. For each slice, B_0 and pre-phase ghost corrections were computed based on the navigator data (Van de Moortele et al., 2002; Thesen et al., 2003). Low-resolution EPI parameters: FOV = 19.2 × 19.2 × 3.9 cm, matrix size = 64 × 64 × 13 (resolution: 3 × 3 × 3 mm), TR = 1000 ms, TE = 30 ms. High-resolution EPI parameters: FOV = 19.2 × 9.6 × 1.56 cm, matrix size = 128 × 64 × 13 (resolution: 1.2 × 1.2 × 1.2 mm), TR = 1000 ms, TE = 30 ms.

Retinotopic mapping

Primary visual cortex (V1) was identified using standard methods (Sereno et al., 1995; DeYoe et al., 1996; Engel et al., 1997). Briefly, to measure the angular component of the retinotopic maps, we used a 45° wedge of flickering checkerboard that rotated slowly (24-s period) about a central fixation point. These stimuli evoked traveling waves of activity across each of the retinotopic visual cortical areas. We fit the measured time series with a sinusoid and calculated the response phase and coherence, separately for each voxel. The phase values measured temporal delay of the fMRI responses relative to the beginning of the experimental cycle and, therefore, corresponded to the polar-angle component of the retinotopic map. V1 was identified as containing a representation of the contralateral hemifield in or near the calcarine sulcus.

Alignment to volume anatomy

After motion correction, the first volume of the first EPI series was aligned directly to a 3D MP-RAGE volume (whole brain) anatomy acquired in a separate scanning session. Automatic image alignment (Nestares and Heeger, 2000) between EPI data and 3D anatomical scans was performed after inverting the voxel intensities in the EPI image to match the T_1 contrast of the inversion-prepared (MP-RAGE) volume anatomy.

Segmentation of cortical surface

The gray/white matter surface was segmented and reconstructed from the MP-RAGE volume anatomy, with SurfRelax (Larsson, 2001). Patches of the occipital cortex in each subject were then inflated and flattened to aid visualization of the localizer data.

Data analysis

To analyze data from the localizer scans, the first 8 time-points (½ cycle) were discarded, and motion compensation was applied to the remaining data from each scan (Nestares and Heeger, 2000). The resulting time series were then high-pass filtered with a cut-off frequency of 2 cycles per scan in the localizer scans and 9 cycles

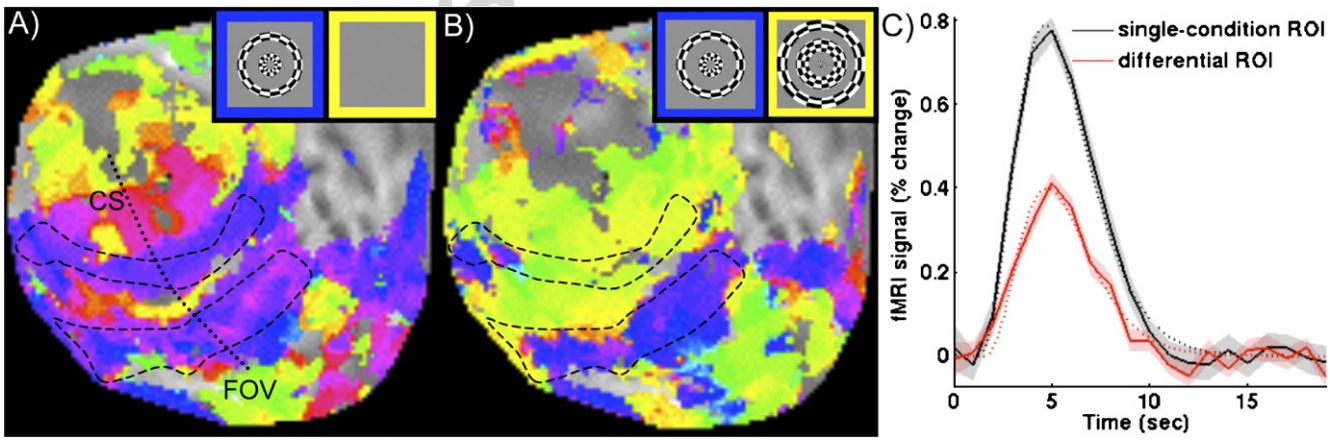

Fig. 2. Localization of visual activity at low resolution (example data for a representative subject). (A) Response phase map for low-resolution single-condition localizer on flattened cortical surface of right hemisphere. Blue/Magenta: activity in phase with the target stimulus. Yellow/Green: activity out of phase with the stimulus. CS: calcarine sulcus (dotted line). FOV: foveal representation. Dashed curves (copied from Fig. 3B): cortical regions responsive to target stimulus from high-resolution differential localizer, measured in a separate scanning session. (B) Phase map for low-resolution differential localizer, visualized on flattened cortical surface. Blue/Magenta: activity in phase with the target stimulus. Yellow/Green: activity in phase with the complement of the target stimulus. (C) Average hemodynamic impulse response functions (HIRFs) for voxels selected by single-condition (black) and differential (red) localizers. Dotted curves: fits to data (time-to-peak: 4.7 s and 4.8 s; peak amplitude: 0.79% and 0.40%).

personal copy

per scan in the HIRF scans (that is, 6 cycles per scan below the stimulus frequency).

For each of the two localizers for each spatial resolution, a region of interest (ROI) was defined as a subregion of primary visual cortex (see following description) that responded strongly to target stimulus. The data were averaged across the two repetitions of the localizer scan, and the time series at each voxel was fit by a sinusoid with the same frequency (10 cycles/scan) as that of the stimulus alternations. For each voxel, we then computed the correlation (technically, coherence) between the measured time series and the best-fit sinusoid (Bandettini et al., 1993; Engel et al., 1997). The coherence is the amplitude of modulation at the stimulus-alternation frequency divided by the integrated amplitude across all frequencies and, hence, was close to 1 if there was a strong stimulus-evoked response with no noise at other frequencies and was close to 0 if there was no stimulus-evoked response or if the response was dominated by noise at other frequencies. If this coherence value exceeded a particular threshold for a localizer scan, the voxel was included in the corresponding ROI. For high-resolution data, the threshold was set to include 500 voxels in the ROI selected by the differential localizer. The same threshold was then used to define the single-condition localizer ROI, which was therefore larger (sizes ranged from 516–951 voxels). For low-resolution data, the threshold was set to include approximately 100 voxels in the differential localizer ROI. Using the same coherence threshold to select voxels for the low resolution, single-condition localizer selected between 159 and 220 voxels.

To estimate the hemodynamic impulse response function (HIRF) in a ROI, the data from each voxel were slice-time corrected, i.e., resampled at the actual acquisition times (using sinc interpolation). Then the data were averaged across voxels in the ROI, and a trial-triggered average was computed by averaging the timecourse following each stimulus presentation across repeated trials. Four repetitions of the HIRF scans were run, with 15 trials per scan, so 60 trials were averaged together for each HIRF estimate. Error bars in the figures indicate the standard error of the mean calculated for 60 trials (i.e., the ROI average for each trial was treated as one observation).

Both single-voxel and ROI average HIRF estimates were fit to a model of the hemodynamic response, which consisted of a difference of gamma functions, using custom Matlab code. The time-to-peak (TTP) and response amplitude (positive or negative) was calculated for the smooth fit to the data. After fitting the estimated HIRF, each voxel was marked as either positively responding, negatively responding or not significantly modulated. A voxel was considered significantly modulated if the absolute value of the peak response was greater than 1.5 times the standard deviation of the time-points 12–20 s after the stimulus presentation. This conservative criterion succeeded in segregating clearly positive and clearly negative responses in all subjects (in all ROIs at both resolutions).

Voxels close to large veins were expected to have high variance and low signal intensity (Duyn, 1995; Lee et al., 1995). Each ROI was therefore divided, in a separate analysis, into two component ROIs: “high-variance” (expected to be associated with large veins) and a remainder (expected to have a smaller contribution from large veins). To select high-variance voxels, a residual time series was computed for each voxel by subtracting the estimated HIRF (trial triggered average) from the time series following each stimulus presentation. The variance of the residual was then normalized by the mean voxel intensity. “High-variance” voxels were identified as the top 20% when voxels in an ROI were ranked by mean-normalized variance. A similar subset of voxels could have been selected simply by ranking the voxels in an ROI by intensity, and selecting the lowest 20%. In a complementary analysis, voxels were also categorized as high- or low-variance by calculating the variance in the localizer scans, and final results were indistinguishable.

Venogram

To verify that mean-normalized variance in functional scans was a sufficient marker of for large veins, we acquired a high-resolution venogram in one subject. The pulse sequence was a 3D FLASH acquisition, with the following parameters: TR = 45 ms; TE = 40 ms; matrix size 384 × 384 × 32; FOV: 220 × 220 × 38.4 cm. Nominal inplane resolution was 0.6 mm and slice (partition) thickness was 1.2 mm. Two scans were averaged together for a total acquisition time of 14 min. Reconstructed images were complex-valued; the phase information was used as a mask to enhance contrast in the magnitude images (Haacke et al., 2004).

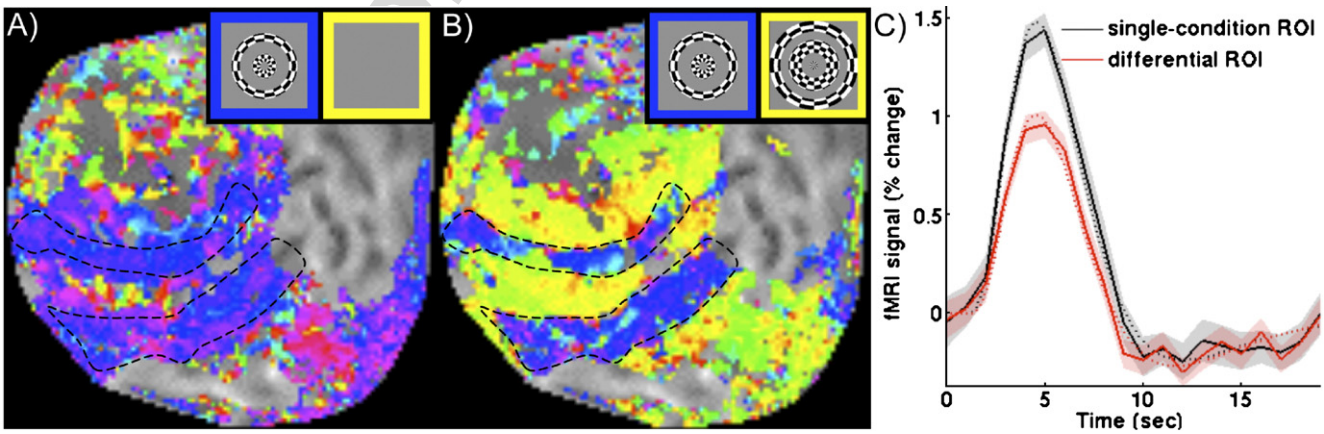

Fig. 3. High-resolution localizers (same subject and same format as Fig. 2). Fits in panel (C): time-to-peak: 4.6 s and 4.6 s; peak amplitude: 1.5% and 1.0%. Because of cortical magnification of the visual field, the representations of the more central annulus are wider than those of the peripheral annulus, even though both subtended 2° of visual angle.

Author's personal copy

Simulation

A model was implemented to simulate the experimental results at high- and low-resolution, for a small region of cortex. The simulated region was 1.2 cm wide and 4.2 cm long, with several large pial veins running across the length of the simulated patch. The simulated volume was 6 mm thick, containing 3 mm of cortex with 1.2 mm of white matter (non-responsive) below and 1.8 mm of CSF above. The central 1.8 cm of the patch was simulated as responding to the target stimulus, flanked on each side by 1.2 cm of cortex responsive to flanking stimuli. At a resolution of 0.15 mm, BOLD responses were estimated for the differential localizer. First, each point in the parenchyma was assigned a response amplitude (1.5%) and phase (0 for the "target" section and π for the "flank" sections). Next, the parenchymal response was blurred with a Gaussian kernel (full width at half-amplitude = 1.5 mm) to simulate blurring by intracortical veins that are spaced, on average, every 1.5–2 mm across the pial surface (Duvernoy et al., 1981). Finally, the response of voxels directly on the pial vein was given the phase of the (blurred) parenchymal response 6 mm "upstream" and a modulation amplitude of 45%. The BOLD response in voxels near the vein was calculated as the complex sum of the response in the parenchyma and the extravascular BOLD effect of the vein, which decreased with the square of the distance from the vein (Ogawa et al., 1993). Finally, the simulated cortical responses were down-sampled to simulate the BOLD response with 1.2- and 3-mm isotropic voxels, as in the high- and low-resolution experiments.

Results

The basic experiment was performed twice in each subject, once with low resolution (3-mm isotropic voxels) and once with high resolution (1.2-mm isotropic voxels). Stimuli are shown in Fig. 1. The experiment consisted of 8 scans: two single-condition localizers (target stimulus vs. blank screen), two differential localizers (target stimulus vs. its complement), and four scans designed to measure the hemodynamic impulse response function, or HIRF (brief presentations of the target stimulus every 20 s).

In both the high- and low-resolution experiments, the regions of interest (ROIs) defined by the differential localizers had a more restricted cortical representation than the single-condition ROIs (Figs. 2 and 3A and B). The amplitude of the average estimated HIRF in the single-condition localizer ROI was, in all subjects, larger than in the differential localizer (Figs. 2C and 3C). Qualitatively, the observed spatial distribution of the HIRF response (data not shown) was better matched to the single-condition localizer than to the differential localizer.

To test for the effects of large veins, which are known to result in signal displacement, each ROI was sorted by mean-normalized variance and divided into two component ROIs (top 20% and remaining 80%). The spatial distributions of the voxels in the high-variance component ROIs were consistent with the expectation that these component ROIs contained many voxels dominated by large-vein effects (Fig. 4A). Voxels marked by high mean-normalized variance were scattered throughout the ROIs and centered on regions in the image with low intensity. Many of these were buried deep in sulci, where large pial veins would be expected, although the high variance in some voxels might be attributed to motion artifacts at edges. (Voxels at image intensity boundaries often have high variance because of their sensitivity to brain motion.) Mean-normalized variance is, however, an indirect indication of the presence of large veins. In one subject, a high-resolution venogram was acquired to verify that voxels marked by high mean-normalized variance contained large veins (Fig. 4B).

In the ROIs defined by the differential localizers, negative HIRF responses were associated with voxels marked as dominated by large veins. The impulse response in the top 20% component ROI, being associated with the signal near large veins, was expected to have a higher amplitude and slightly longer latency than the average response in the rest of the ROI. This was observed, in all subjects, in both of the high-resolution ROIs (single-condition and differential localizer, Figs. 5C and D) and in the single-condition ROI for the low-resolution data (Fig. 5A). For the differential ROI at low resolution, however, the average response in the large-vein weighted (top 20% mean-normalized variance) component ROI was negative (Fig. 5B) in 3 of 4

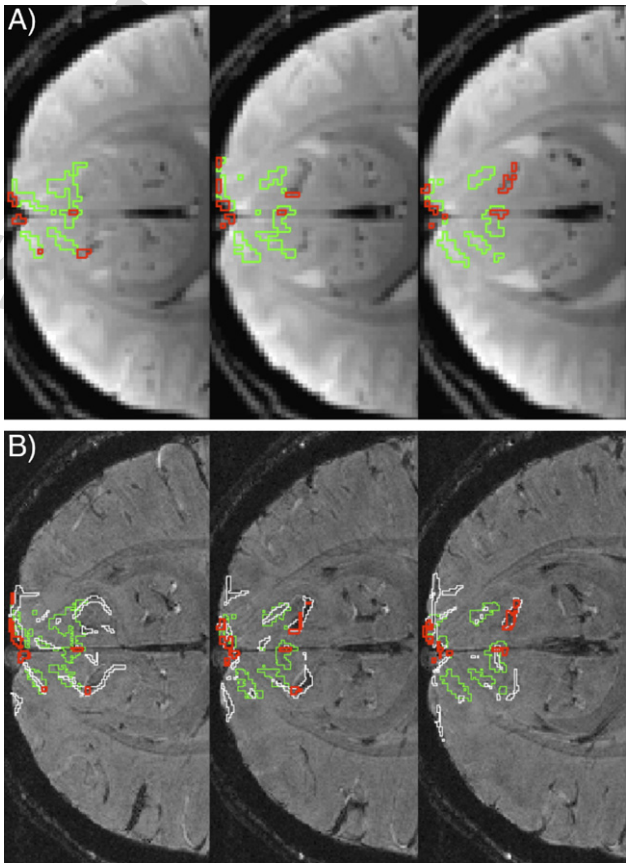

Fig. 4. Spatial distribution of high-variance voxels, for a representative subject. (A) Grayscale image: mean EPI image from one functional run. Overlay: ROI defined by high-resolution, single-condition localizer. High-variance voxels, shown in red, lie predominantly in regions with short T_2^* that appear dark (low intensity) in the images. Voxels in the remaining 80% of the original ROI are shown in green. (B) Grayscale image: venogram (in-plane resolution 0.6 × 0.6 mm). Overlay: identified veins (white) in primary visual cortex, high-variance (red) and remainder (green) ROIs from (A). Data were acquired on different days with different slice placement/orientation, so pattern of high-variance voxels appears slightly different in the two panels, due to resampling. In the high-variance component ROI, 48% of the voxels were aligned with image regions marked as veins in the venogram; 15% of the remainder ROI were marked as veins.

Author's personal copy

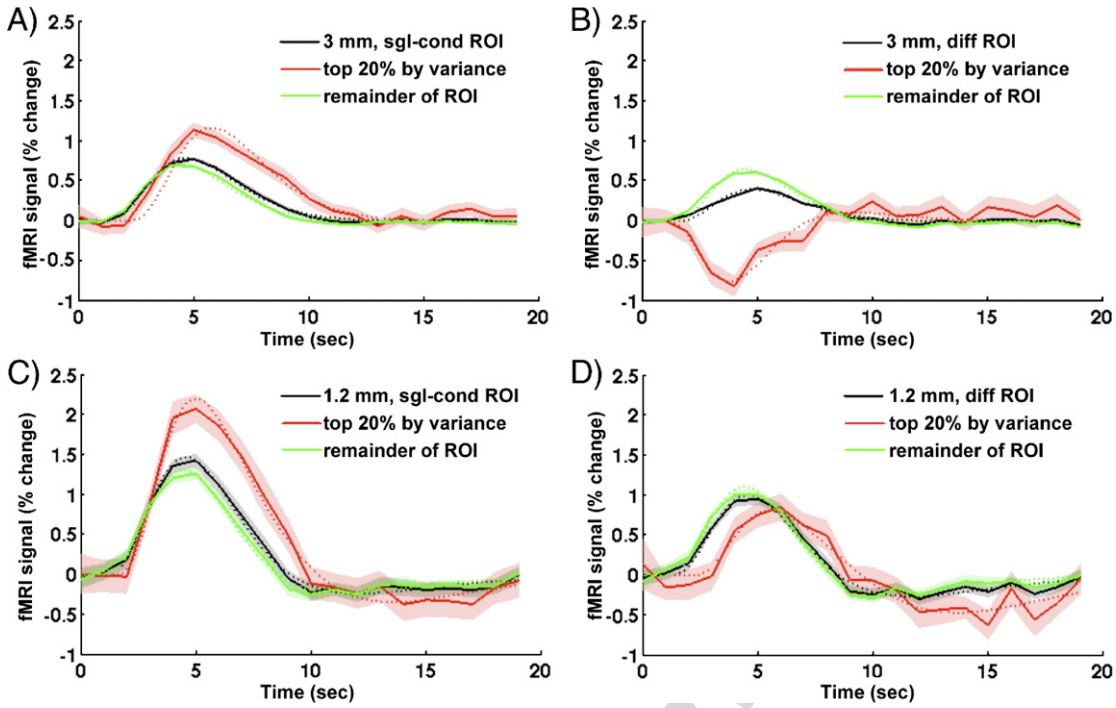

Fig. 5. HIRF estimates (example data for representative subject). The ROI defined by each localizer was divided into two sub-populations: the top 20% when sorted by residual variance normalized by mean intensity, and the remainder. Black curves, HIRF estimates including all voxels. Red, HIRF estimates from high-variance voxels. Green, HIRF estimates from the remaining low-variance voxels. (A) Low-resolution data, ROI defined by single-condition localizer. (B) Low resolution, differential localizer. (C) High resolution, single-condition localizer. (D) High resolution, differential localizer. In the single-condition ROIs, high-variance component ROIs have large amplitude and long latency responses, characteristic of large draining veins. High-variance component ROI responds negatively in the low-resolution, differential ROI (panel B, red curve).

subjects. (In the fourth subject, the average response was markedly delayed with low amplitude). Negative hemodynamic responses have been reported previously in cortical territory or voxels that flank stimulated cortex (Shmuel et al., 2002; Devor et al., 2005), but not in voxels that correspond directly to the V1 cortical representation of a visual stimulus. The negative responses that we measured were in voxels selected by the differential localizer and were reliably located within the cortical representation of the stimulus. Fitting the HIRF on a voxel-by-voxel basis revealed that, across subjects, 13±4% (mean±SD) of the voxels selected by the low-resolution differential localizer responded negatively to the target stimulus in the HIRF scans. In contrast, few of the voxels selected by the high-resolution differential localizer responded negatively in the HIRF scans (4.2±1.2%). When the low-resolution ROIs were divided into high-variance and remainder ROIs, between 20% and 65% (39.6±19.4%) of the voxels in the high-variance component of the differential ROI had negative HIRFs. The fact that voxels with negative HIRFs comprised a higher percentage of the high-variance ROI than the remainder suggests significant overlap between the population of voxels with high mean-normalized variance and the population with negative HIRFs. In the high-resolution data, a smaller fraction of the high-variance ROIs had negative HIRFs (6.3±0.9%).

Similar results were revealed when we calculated, on a voxel-by-voxel basis, whether a positive response to one localizer predicted a positive or a negative response to the other localizer (Fig. 6). A positive response to the single-condition localizer did not predict a positive response to the differential localizer (Figs. 6A and C). This is consistent both with the results shown in Figs. 2 and 3 and with other studies that have shown differential protocols do a better job of limiting the ROI to the cortical representation of the stimulus. Voxels outside the cortical representation of the stimulus, selected by the single-condition localizer, should respond 180° out of phase with (negatively to) the target stimulus in the differential localizer. However, if the differential localizer selects voxels only within the cortical representation of the target stimulus, then all voxels in the differential ROI should respond in phase with the stimulus in the single-condition localizer. This is not what we observed. In the low-resolution data, a large number of voxels in the differential localizer ROI responded negatively to the single-condition localizer (17±5%). Few voxels exhibited this behavior in the high-resolution experiment (1.6±0.9%), and the few that did were contained in the high-variance component ROI (Fig. 6D). This result was independent of the method of ROI definition: the results were similar whether a stringent or liberal coherence threshold was used to define the ROIs, and whether the ROIs were matched for number of voxels or for coherence threshold.

After the high-variance voxels were removed from the low-resolution differential ROIs, the amplitude of the average HIRF response in the remainder ROIs was larger (yellow bar above blue in Fig. 7A). This finding suggested that the relatively low mean response observed when averaging across all voxels in the differential ROI was due to inclusion of unresponsive or negatively responding voxels (negative red bar in Fig. 7A). Consistent with the expectation for large-vein weighted ROIs, latencies in the high-variance component ROIs were on average 8% longer (in all but the low-resolution differential ROI), and amplitudes were 50%

Author's personal

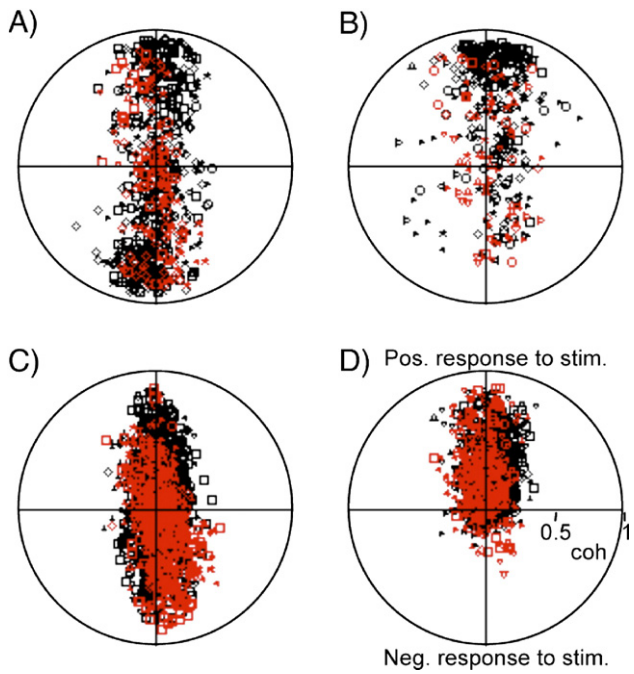

Fig. 6. Amplitude and phase of responses to one type of localizer for voxels in ROIs selected by the other type of localizer (combined across subjects, each of which is represented by a different symbol shape). (A) Low-resolution experiment, responses to differential localizer in voxels selected by single-condition localizer. Distance from origin indicates coherence. Azimuth indicates response phase (phase of stimulus is along positive *y* axis, such that voxels responding negatively to stimulus are plotted along the negative *y* axis). Each plot symbol represents a voxel (red: voxels in high-variance component ROI; black: remainder ROI). The tendency toward longer latency (Figs. 4 and 7) is seen in the slight counter-clockwise rotation of the distribution in the high-variance voxels. (B) Low-resolution, responses to single-condition localizer in voxels belonging to differential ROI. Negative responses (on negative *y* axis) are not expected if differential localizer is more accurate than single-condition. (C) High resolution, responses to differential localizer in voxels belonging to single-condition localizer ROI. (D) High resolution, responses to single-condition localizer in voxels selected by differential localizer. Only in the high-resolution data did the differential localizer consistently predict responses to the single-condition localizer.

larger than in the low-variance component ROIs (Fig. 7). Although the high-resolution localizers showed less of a discrepancy between the two methods of voxel selection (single-condition and differential localizers), there was still some evidence of contamination by large-vein effects (red bars above yellow in Fig. 7).

An illustration of the association between negative hemodynamic responses and voxels that have high mean-normalized variance is shown in Fig. 8, which summarizes the primary finding of this paper. The differential localizers, at high and low resolution, appeared to do a better job of selecting voxels that matched the known cortical organization of the stimulus response. In general, this ROI was contained within the ROI selected by the single-condition localizer. In the low-resolution experiment, however, an average (across subjects) of 13% of the voxels in the ROI defined by the differential localizer responded negatively to the target stimulus when it was presented in a single-condition protocol (red outline in Figs. 8A, B and red curve in Fig. 8C). These inconsistent responses, voxels that responded positively to the stimulus in the differential localizer but negatively in the single-condition localizer or impulse response measurement, were not expected inside the cortical representation of the stimulus. They were associated with voxels that had high mean-normalized variance, indicative of the presence of a large vein that dominated the local BOLD response when a single-condition protocol was used. Similar inconsistent responses were observed in the high-resolution experiment, but a much smaller proportion of voxels showed these effects.

We implemented a model to simulate the experimental results (Fig. 9). Large pial veins were modeled to displace the BOLD signal 6 mm across the boundary between flanker and target cortical territories. The volume of voxels influenced by the pial vein at high and low resolution was compared by quantifying how many voxels within the target cortical territory were modulated out of phase with the target stimulus. For a range of reasonable parameters, we found that the volume influenced by the vein was 2–3 times larger in the low-resolution simulation than the high resolution. The exact volume depended on the magnitude of signal displacement as well as the relative magnitude of the BOLD response in the large veins, the extent of the extravascular effects, and the density of the pial vessels. While this model succeeded in demonstrating the advantage of high resolution for minimizing large-vein effects, it failed to explain the source of the negative BOLD response in the single-condition stimulus. It also predicted, inconsistent with the experimental data, that BOLD modulation should be absent in only a small section of the pial vessel in the differential protocol.

Discussion

Several of the findings of this paper confirm known characteristics of BOLD fMRI: (1) voxels containing large veins can be identified by high variance relative to mean intensity and have slower hemodynamic impulse response functions, (2) excluding voxels dominated by large veins from ROIs can improve the accuracy of subsequent analysis, and (3) differential experimental designs provide more accurate cortical maps than single-condition designs. What was not expected was the observation that many voxels identified by differential localizers responded negatively to the stimulus when it was presented in a single-condition protocol. Consequently, differential ROIs were not simply a subset of the single-condition ROIs. This spatial mismatch between localizers resulted in poor estimates of stimulus-evoked response in low-resolution data when a differential localizer was used to analyze data presented in a single-condition protocol.

The best interpretation of our data is that a strong negative BOLD signal is displaced by large veins into the target cortical territory in single-condition (and impulse response) protocols, but not in differential protocols. The negative responses that we observed may have the same origin as what have been called "negative BOLD" responses. In negative BOLD, the hemodynamic response in cortex flanking the stimulated cortex is consistently negative, either because of neural suppression or hemodynamic "stealing" (Shmuel et al., 2002; Devor et al., 2005). In our data, however, the negative responses were not observed only in flanking cortex. The negative responses were also observed in many voxels that appeared to be well within the cortical representation of the stimulus and that responded positively to the stimulus in a differential protocol. Presumably because of their proximity to large veins (marked by high mean-normalized variance in the HIRF scans), these voxels responded negatively

Author's personal copy

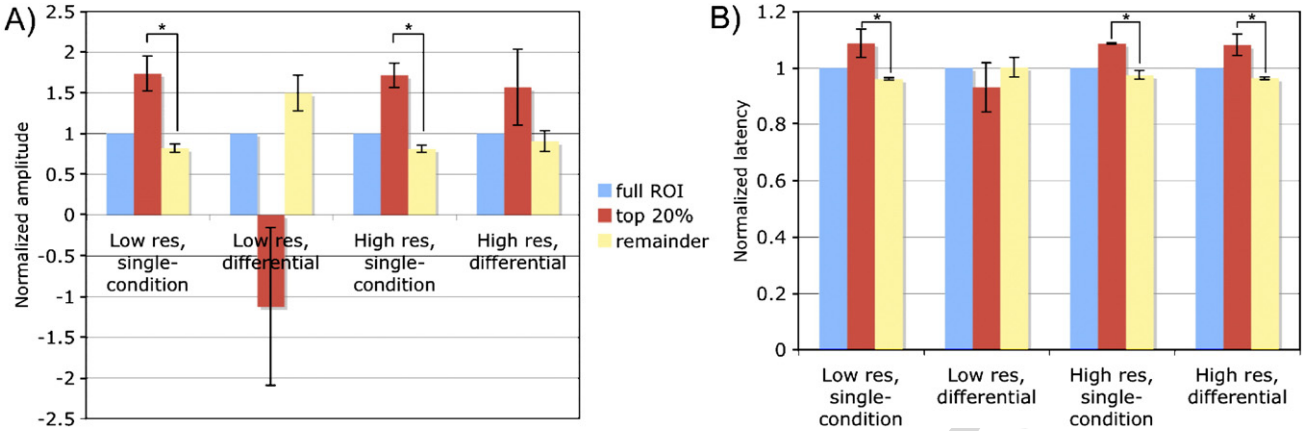

Fig. 7. Summary of response amplitudes and latencies in all ROIs and all subjects. (A) Peak amplitude in component ROIs (red: top 20% by variance; yellow: remainder), normalized by peak amplitude in full ROI (blue) and then averaged across subjects. Error bars reflect standard error of the mean across subjects (N=4). (B) Time to peak for component ROIs (normalized and averaged as in A). In all but the low-resolution differential ROIs, the high-variance voxels had higher amplitude responses and longer latencies than the remainder of the voxels (asterisk: $p<0.05$ in one-tailed paired t-test with 3 degrees of freedom).

to a single-condition stimulus, even though they occupied the cortical representation of the stimulus.

We conclude that the negative BOLD signal is being displaced, rather than blurred, because we measured inconsistent responses in voxels *within* the cortical representation of a stimulus, while a blurring or point-spread function model (Engel et al., 1997; Parkes et al., 2005) would predict inconsistent responses *outside* of the cortical representation. The key difference is measured in voxels that are within the cortical representation of the stimulus, but which contain (or are dominated by) a large vein that collects blood from outside the cortical representation of the stimulus. In a blurring model, these voxels would respond in phase with the target stimulus. Only displacement of signal from outside the target cortical territory would explain the strong out-of-phase responses we observed within the cortical representation of the stimulus.

We found that acquiring data with smaller voxels (higher resolution) reduced this problem. As illustrated in Fig. 9, BOLD effects of large veins should dominate a smaller fraction of the voxels at higher resolution. When the vascular architecture is better resolved, some voxels contain only small veins, and some contain almost exclusively large vessels, instead of every voxel having a mixture of large and small vein effects. In the simulation results shown, high resolution reduced the volume affected by the vein by a factor between 2 and 3. The density of pial veins, and their location in relation to the cortical boundary, will determine what percentage of the target ROI these volumes represent. With one pial vessel every 3 mm, the large vein-biased voxels represented 17% and 7% of the low- and high-resolution synthetic data, in good agreement with the experimental data.

One caveat to be offered about the data presented is the method of selection for "large vessel" voxels. We based our identification of these voxels on residual variance. Any voxel for which we had a poor model of the impulse response, whether this was due to motion or some other inconsistency in response throughout the four scans that contributed to the trial-triggered average, would have been marked as having high variance relative to its mean. This can explain, for example, the inclusion of some voxels at the boundary between cortex and CSF in the "top 20%" ROIs. This

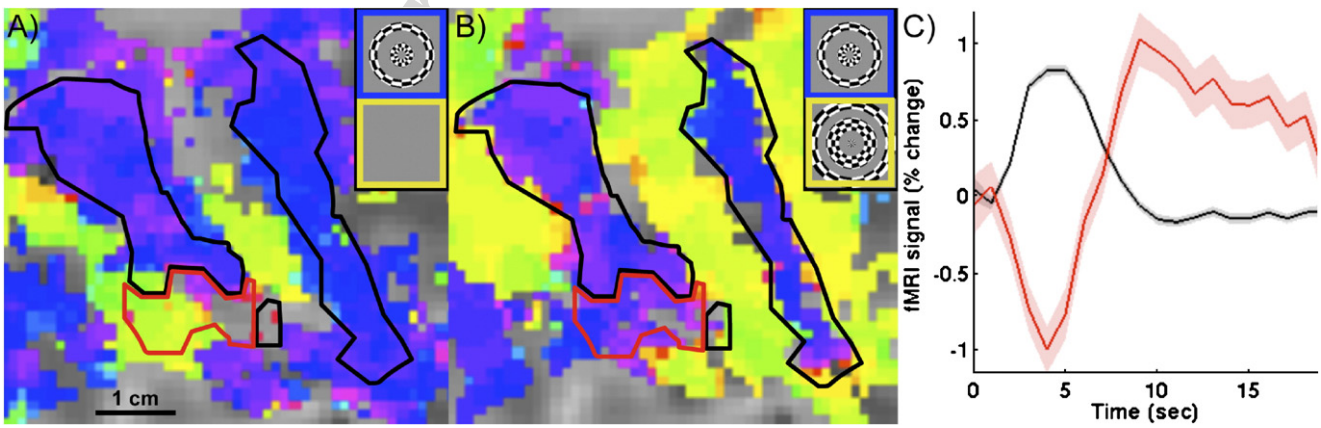

Fig. 8. Association between negative responses to briefly presented stimulus and voxels with high mean-normalized variance. (A) Responses to low-resolution single-condition localizer, as in Fig. 2 but for a different subject. Overlaid are component ROIs defined by the differential localizer. Red, top 20% of voxels by mean-normalized variance. Black, remainder. (B) Responses to low-resolution differential localizer, showing a pattern that is more restricted and better matched to known cortical organization. Same ROI overlay. (C) Estimated HIRFs for component ROIs, showing a strongly negative response in the putative large-vein voxels (red trace, red ROI in A), and a positive response in the residual ROI (black trace).

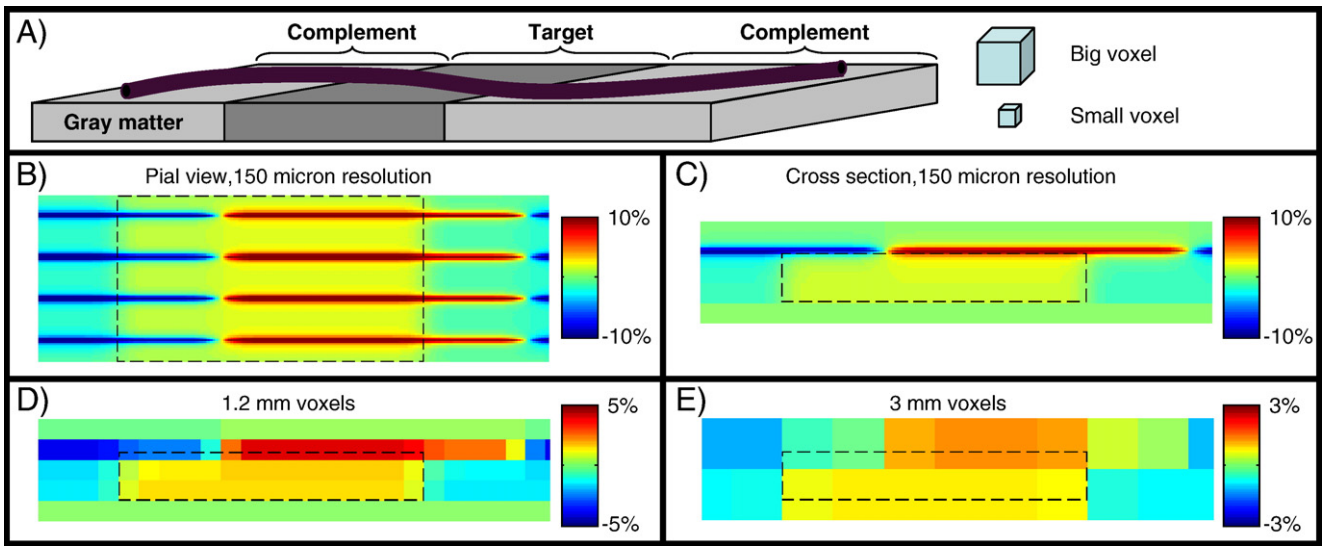

Fig. 9. Model simulation. (A) The target stimulus and its complement stimulated different regions of cortex in retinotopic visual cortex. A large draining pial vein traversing both regions is shown in purple. (B) Spatial distribution of responses in parenchyma (capillaries and intracortical veins) and large pial vessels (separated by 3 mm), simulated with 150-μm resolution, viewed from the pial surface. Blue, portions in which BOLD signal increased in response to the presentation of the flanking stimulus (or decreased in response to the target stimulus). Reds and yellows, responses were in phase with the target stimulus. Green, unmodulated regions. Signal was displaced in the large pial vein, 6 mm in this example. Therefore, voxels near the boundary of the stimulus representation have conflicting contributions from the parenchyma and the vein. The magnitude of the contribution of the vein depends on orientation and diameter and decays as $1/r^2$ (r is the distance from the vein); for this example, the maximum BOLD modulation in the vein was simulated as being 45% in the vein; 1.5% in parenchyma. The parenchymal response was blurred with a Gaussian kernel with full width at half-maximum of 1.5 mm to represent the averaging effects of intracortical veins. Black dashed box indicates region in which underlying parenchymal response is in phase with target presentation. (C) Cross-sectional view of the simulation. Unmodulated regions (white matter and CSF) shown in green. (D) Simulation of the high-resolution experiment, generated by down-sampling from 150 μm to 1.2 mm. In the "true" target region, 7% of the voxels responded negatively to the target stimulus because of the presence of the vein. (E) Simulated low-resolution data. Because of partial volume effects, the fMRI responses were weaker. At low resolution, 17% of the voxels in the simulated target region responded negatively to the target stimulus.

method for voxel selection was, however, validated by the alignment of these voxels with identified veins in a high-resolution venogram, and by the temporally delayed and large (or negative) amplitude responses of these voxels. Our results are consistent with previous studies using mean-normalized variance to select vein-weighted voxels (de Zwart et al., 2005).

Our results, showing a significant effect from displacement of the BOLD signal, are also consistent with previous studies that attempted direct measurement of the effective localization accuracy of BOLD-based fMRI measurements. A study by Disbrow et al. (2000) compared, in the same monkeys, the centroids of electrophysiological and fMRI maps evoked by hand and face stimulation. They found that the greatest discrepancies between their single-condition fMRI maps and the electrode recording maps occurred along an axis that was parallel to the cortical surface, even though this dimension (in their case) had the highest imaging resolution. These localization errors are consistent with the displacement of BOLD response by large pial veins. Disbrow et al. (2000) found discrepancies between electrophysiological and fMRI centroids as large as 1 cm, comparable to the largest displacement we measured. Tjandra et al. (2005) compared the center of mass of BOLD activation against cerebral blood flow (CBF), a technique known to be more spatially reliable. The BOLD centroids were displaced toward large veins visible in T_2^*-weighted venograms. While some theoretical calculations predict that the displacement of BOLD response in large veins is limited to a few millimeters because of dilution effects (Turner, 2002), there is also direct evidence from optical imaging that large-vein responses can extend for more than a centimeter, even with focal neural activity (Shoham and Grinvald, 2001).

An important special case of the differential protocol is the "traveling wave" protocol frequently used in retinotopic mapping (Engel et al., 1994). When the stimulus moves continuously through visual space, adjacent regions of cortex are stimulated sequentially, largely eliminating modulation of deoxyhemoglobin changes in large veins. Our data suggest that for voxels near large veins, which can be as many as 20% of the voxels, a discrepancy should be expected between the visual field representation mapped with a traveling wave protocol and the responses to single-condition stimulus protocols (e.g., a briefly flashed stimulus presentation).

The use of a localizer scan to select voxels for analysis is a common practice, particularly in visual neuroscience. A differential experimental protocol reduces the contribution of large veins, and the improved spatial accuracy is expected to yield improved localization of visual responses. However, in our low-resolution data, we found that ROIs defined by differential localizers contained a number of voxels that had a negative hemodynamic response to a stimulus that was subsequently presented (briefly, or in a single-condition protocol) at the same location in the visual field. Therefore, even though spatial accuracy was apparently good, the functional representation was confounded by displaced negative hemodynamic responses. Identifying and removing voxels dominated by large-vein responses ameliorated the problem. We have interpreted this result in the context of a BOLD response that is displaced, not isotropically blurred, on the cortical surface. Our

Author's personal

finding highlights an important aspect of fMRI with GE BOLD: the spatial errors introduced by large-vein effects are significant, and they are unpredictable because they depend on the local geometry of the vasculature. These errors are reduced at higher spatial resolutions and can be eliminated by identification and removal of voxels dominated by large veins.

Acknowledgments

The authors would like to thank Justin Gardner for his assistance with data analysis and comments on the manuscript, and Shani Offen and Jennifer Schumacher for assistance with data collection. This project was funded by awards from the McKnight Foundation and the Seaver Foundation.

References

Bandettini, P.A., Jesmanowicz, A., Wong, E.C., Hyde, J.S., 1993. Processing strategies for time-course data sets in functional MRI of the human brain. Magn. Reson. Med. 30 (2), 161–173.

Bonhoeffer, T., Grinvald, A., 1993. The layout of iso-orientation domains in area 18 of cat visual cortex: optical imaging reveals a pinwheel-like organization. J. Neurosci. 13, 4157–4180.

Cheng, K., Wagooner, R.A., Tanaka, K., 2001. Human ocular dominance columns as revealed by high-field functional magnetic resonance imaging. Neuron 32, 359–374.

Devor, A., Ulbert, I., Dunn, A.K., Narayanan, S.N., Jones, S.R., Andermann, M.L., Boas, D.A., Dale, A.M., 2005. Coupling of the cortical hemodynamic response to cortical and thalamic neuronal activity. Proc. Natl. Acad. Sci. 102 (10), 3822–3827.

DeYoe, E.A., Carman, G.J., Bandettini, P.A., Glickman, S., Wieser, J., Cox, R.W., Miller, D., Neitz, J., 1996. Mapping striate and extrastriate visual areas in human cerebral cortex. Proc. Natl. Acad. Sci. U. S. A. 93, 2382–2386.

de Zwart, J.A., Silva, A.C., van Gelderen, P., Kellman, P., Fukunaga, M., Chu, R., Koretsky, A.P., Frank, J.A., Duyn, J.H., 2005. Temporal dynamics of the BOLD fMRI impulse response. NeuroImage 54, 667–677.

Disbrow, E.A., Slutsky, D.A., Roberts, T., Krubitzer, L.A., 2000. Functional MRI at 1.5 Tesla: a comparison of the blood oxygenation level-dependent signal and electrophysiology. Proc. Natl. Acad. Sci. U. S. A. 97 (17), 9718–9723.

Duvernoy, H.M., Delon, S., Vannson, J.L., 1981. Cortical blood vessels of the human brain. Brain Res. Bull. 7, 519–579.

Duyn, J.H., 1995. Effects of large vessels in functional magnetic resonance imaging at 1.5 T. Int. J. Imaging Sys. Technol. 6, 245–252.

Engel, S.A., Rumelhart, D.E., Wandell, B.A., Lee, A.T., Glover, G.H., Chichilnisky, E.M., Shadlen, M.N., 1994. fMRI of human visual cortex. Nature 369, 525.

Engel, S.A., Glover, G.H., Wandell, B.A., 1997. Retinotopic organization in human visual cortex and the spatial precision of functional MRI. Cereb. Cortex 7, 181–192.

Farzaneh, F., Riederer, S.J., Pelc, N.J., 1990. Analysis of T2 limitations and off-resonance effects on spatial resolution and artifacts in echo-planar imaging. Magn. Reson. Med. 14, 123–139.

Fleysher, L., Fleysher, R., Heeger, D.J., Inati, S., 2005. High resolution fMRI using a 3D multi-shot EPI sequence. Proc. Int. Soc. Magn. Reson. Med. 13, 2685.

Grinvald, A., Lieke, E.E., Frostig, R.D., Hildesheim, R., 1994. Cortical point-spread function and long-range lateral interactions revealed by real-time optical imaging of macaque monkey primary visual cortex. J. Neurosci. 14 (5), 2545–2568.

Haacke, E.M., Xu, Y., Cheng, Y.-C.N., Reichenbach, J.R., 2004. Susceptibility weighted imaging (SWI). Magn. Reson. Med. 52, 612–618.

Heeger, D.J., Ress, D., 2002. What does fMRI tell us about neuronal activity? Nature 3, 142–151.

Larsson, J., 2001. Imaging vision: functional mapping of intermediate visual processes in man. Stockholm, Sweden, Karolinska Institutet. PhD thesis.

Lee, A.T., Glover, G.H., Meyer, C.H., 1995. Discrimination of large venous vessels in time-course spiral blood-oxygen-level-dependent magnetic-resonance functional neuroimaging. Magn. Reson. Med. 33, 745–754.

Nestares, O., Heeger, D.J., 2000. Robust multiresolution alignment of MRI brain volumes. Magn. Reson. Med. 43, 705–715.

Norris, D.G., Zysset, S., Mildner, T., Wiggins, C.J., 2002. An investigation of the value of spin-echo-based fMRI using a Stroop color-word matching task and EPI at 3 T. NeuroImage 15 (3), 719–726.

Ogawa, S., Menon, R.S., Tank, D.W., Kim, S.-G., Merkle, H., Ellerman, J.M., Ugurbil, K., 1993. Functional brain mapping by blood oxygenation level-dependent contrast magnetic resonance imaging. A comparison of signal characteristics with a biophysical model. Biophys. J. 64 (3), 803–912.

Parkes, L.M., Schwarzbach, J.V., Bouts, A.A., Deckers, R.H.R., Pullens, P., Kerskens, C.M., Norris, D.G., 2005. Quantifying the spatial resolution of the gradient echo and spin echo BOLD response at 3 Tesla. Magn. Reson. Med. 54, 1465–1472.

Pfeuffer, J., Van de Moortele, P.F., Yacoub, E., Shmuel, A., Adriany, G., Andersen, P., Merkle, H., Garwood, M., Ugurbil, K., Hu, X., 2002. Zoomed functional imaging in the human brain at 7 Tesla with simultaneous high spatial and high temporal resolution. NeuroImage 17 (1), 272–286.

Pruessmann, K.P., 2004. Parallel imaging at high field strength: synergies and joint potential. Top. Magn. Reson. Imaging 15, 237–244.

Sereno, M.I., Dale, A.M., Repas, J.B., Kwong, K.K., Belliveau, J.W., Brady, T.J., Rosen, B.R., Tootell, R.B.H., 1995. Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. Science 268 (5212), 889–893.

Shmuel, A., Yacoub, E., Pfeuffer, J., Van de Moortele, P.-F., Adriany, G., Hu, X., Ugurbil, K., 2002. Sustained negative BOLD, blood flow and oxygen consumption response and its coupling to the positive response in the human brain. Neuron 36, 1195–1210.

Shoham, D., Grinvald, A., 2001. The cortical representation of the hand in macaque and human area S-I: high resolution optical imaging. J. Neurosci. 21 (17), 6820–6835.

Thesen, S., Kruger, G., Muller, E., 2003. Absolute correction of B0 fluctuations in echo-planar imaging. Proc. Int. Soc. Magn. Reson. Med. 11, 1025.

Tjandra, T., Brooks, J.C.W., Figueiredo, P., Wise, R., Matthews, P.M., Tracey, I., 2005. Quantitative assessment of the reproducibility of functional activation measured with BOLD and MR perfusion imaging: implications for clinical trial design. NeuroImage 27, 393–401.

Turner, R., 2002. How much cortex can a vein drain? Downstream dilution of activation-related cerebral blood oxygenation changes. NeuroImage 16, 1062–1067.

Van de Moortele, P.F., Pfeuffer, J., Glover, G.H., Ugurbil, K., Hu, X., 2002. Respiration-induced B0 fluctuations and their spatial distribution in the human brain at 7 Tesla. Magn. Reson. Med. 47 (5), 888–895.

Yacoub, E., Duong, T.Q., Van de Moortele, P.-F., Lindquist, M., Adriany, G., Kim, S.-G., Ugurbil, K., Hu, X., 2003. Spin-echo fMRI in humans using high spatial resolutions and high magnetic fields. Magn. Reson. Med. 49, 655–664.

Yacoub, E., Van de Moortele, P.F., Shmuel, A., Ugurbil, K., 2005. Signal and noise characteristics of Hahn SE and GE BOLD fMRI at 7 T in humans. NeuroImage 24 (3), 738–750.

Author's personal copy