Whole-brain, time-locked activation with simple tasks revealed using massive averaging and model-free analysis

Javier Gonzalez-Castillo[a,1], Ziad S. Saad[b], Daniel A. Handwerker[a], Souheil J. Inati[c], Noah Brenowitz[a], and Peter A. Bandettini[a,c]

[a]Section on Functional Imaging Methods, Laboratory of Brain and Cognition, [b]Scientific and Statistical Computing Core, and [c]Functional MRI Facility, National Institute of Mental Health, National Institutes of Health, Bethesda, MD 20892

Edited by Robert Desimone, Massachusetts Institute of Technology, Cambridge, MA, and approved February 21, 2012 (received for review December 21, 2011)

The brain is the body's largest energy consumer, even in the absence of demanding tasks. Electrophysiologists report on-going neuronal firing during stimulation or task in regions beyond those of primary relationship to the perturbation. Although the biological origin of consciousness remains elusive, it is argued that it emerges from complex, continuous whole-brain neuronal collaboration. Despite converging evidence suggesting the whole brain is continuously working and adapting to anticipate and actuate in response to the environment, over the last 20 y, task-based functional MRI (fMRI) have emphasized a localizationist view of brain function, with fMRI showing only a handful of activated regions in response to task/stimulation. Here, we challenge that view with evidence that under optimal noise conditions, fMRI activations extend well beyond areas of primary relationship to the task; and blood-oxygen level-dependent signal changes correlated with task-timing appear in over 95% of the brain for a simple visual stimulation plus attention control task. Moreover, we show that response shape varies substantially across regions, and that whole-brain parcellations based on those differences produce distributed clusters that are anatomically and functionally meaningful, symmetrical across hemispheres, and reproducible across subjects. These findings highlight the exquisite detail lying in fMRI signals beyond what is normally examined, and emphasize both the pervasiveness of false negatives, and how the sparseness of fMRI maps is not a result of localized brain function, but a consequence of high noise and overly strict predictive response models.

fMRI | activation extent | transient responses | clustering

For years, positive gamma-like responses were the primary type of blood-oxygen level-dependent (BOLD) response that researchers used as indirect markers of neuronal activity. Other methods for extracting neural information from BOLD time series are increasingly popular today. Spontaneous fluctuations recorded in the absence of a controlled stimulus have rapidly gained attention and shown great potential in the study of normal (1, 2) and abnormal brain function (3, 4). Multivariate methods have demonstrated that detailed information about stimulus input can be obtained by jointly analyzing activity in voxels that show no significance using conventional univariate analysis techniques. Even within the framework of univariate analysis, consideration of BOLD responses other than conventional positively sustained responses, such as negatively correlated or stimulus onset/offset responses, has proven useful at differentiating auditory or visual stimuli within primary sensory cortices (5, 6).

Despite all such evidence highlighting the exquisite sensitivity of the BOLD contrast to underlying brain function, few block-design task-based functional MRI (fMRI) studies examine temporal responses other than the conventional positively sustained gamma-like response or conduct secondary analysis on areas of no statistical significance. Factors contributing to this practice may include the lack of desire to look at temporal dynamics once a region is labeled statistically significant, or the fact that conventional responses have proven sufficient to uncover the neuronal correlates of a myriad of human behaviors. Unfortunately, if, as the previous discussion suggests, true neuronal responses are continuously passing undetected in fMRI, our conceptualizations of brain function based on task-based fMRI research might be incomplete.

As Lieberman and Cunningham stated previously (7), a long-standing preoccupation with the reduction of false-positives in fMRI creates a bias toward reporting only large and obvious effects, neglecting what perhaps represents more subtle complex cognitive and affective processes. Here, we explore this hypothesis in detail and evaluate whether the sparseness of task-based fMRI activation maps is real or a consequence of noise levels and modeling decisions. We approach this question using low-noise fMRI time-series generated by combining unconventionally large amounts of data (100 runs per subject). With these data, we also evaluate how regional differences in BOLD response may reveal how distant regions collaborate during a particular task.

What Is the True Extent of BOLD Activations? Previous research has shown that if fMRI noise is reduced by time-series averaging, activation area significantly increases with number of averaged runs (8, 9). Fast increases in activation area during initial averaging stages were followed by a progressive decrease in the rate of area growth with averaging. Still, no asymptotic behavior was reported. Moreover, voxels with subtle hemodynamic responses not strong enough to attain significance with fewer trials showed no significant differences in hemodynamic delay from voxels that were significantly active with fewer trials (8). This finding implies that increases in activation area could not be accounted for as being the result of unaccounted-for hemodynamic delay differences (i.e., large vessels). Similarly, the use of more versatile response models (10) or modeling of additional hemodynamic response shapes (e.g., negatively correlated or stimulus onset/offset responses) have been reported to also increase activation area (5, 6) and account for additional variability (11).

Understanding how noise levels and choice of a predictive BOLD-response model influences fMRI activation maps is a necessary step toward comprehending how much procedural decisions shape fMRI results and obscure the "true" amount of neuronal resources recruited by an experimental perturbation. In other words, when subjects perform a task, is only a reduced set of isolated regions actively recruited? Or, does the majority of the brain show BOLD signal modulation consistent with task timing but pass undetected? To answer these questions we acquired

Author contributions: J.G.-C. and P.A.B. designed research; J.G.-C. performed research; Z.S.S. contributed new reagents/analytic tools; J.G.-C., Z.S.S., D.A.H., S.J.I., and N.B. analyzed data; and J.G.-C., D.A.H., and P.A.B. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Data deposition: The data reported in this paper have been deposited in the Xnat Central database, https://central.xnat.org/ (username: gonzalezcastilloj).

[1]To whom correspondence should be addressed. E-mail: javier.gonzalez-castillo@nih.gov.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1121049109/-/DCSupplemental.

a total of 100 functional runs (i.e., 500 trials) in each of three subjects as they were engaged on a visual stimulation plus letter/number discrimination task. A block paradigm with alternating active (20 s) and resting (40 s) epochs was used (*SI Methods*). We first calculated a single average time-series per subject using all 100 available runs. Averaging reduced noise floor by a factor of six and allowed observation of a variety of BOLD responses time-locked with the task across the whole brain. To evaluate effects on activation extent, we subsequently computed activation maps using incremental amounts of data (N_{runs}), where N_{runs} was varied between 1 and 100. To minimize any bias, we used multiple run permutations at each N_{runs} level, and computed maps of significant activation ($P_{Bonf} < 0.05$ and $P_{FDR} < 0.05$) using three different models (Fig. 1):

(*i*) a sustained response model (SUS) consisting of the convolution of a gamma-variate function with a boxcar function that follows the experimental paradigm;
(*ii*) an onset+sustained+offset model (OSO) that includes transitory responses at blocks onset/offsets in addition to the sustained response (6); and
(*iii*) an unconstrained model (UNC) consisting of 30 impulse functions spanning the duration of a single on/off cycle (60 s) and therefore setting no a priori constraints on response shape other than agreement with task periodicity.

These models were selected because of their increasing flexibility in fitting hemodynamic responses of variable shape. It is not the purpose of this work to extensively evaluate the performance of predictive-response models, but to see how additional flexibility affects activation maps. For detailed model comparisons and other models, please refer to refs. 10, 11, 13, and 14.

Data indicate that activation area varies substantially with N_{runs} and response model, raising important questions about how to interpret fMRI activation maps or the meaningfulness of statistical thresholds.

Do Regional Differences in BOLD Response Hold Information About Brain Function? Hemodynamic responses vary across subjects and regions (15, 16). This variability has been attributed in part to differences in underlying neuronal processing roles (17). Estimation of the time course of the BOLD response associated with short-duration activity is common, and interregional differences in its temporal characteristics are often presented as evidence to postulate putative differences in the role that such regions play with respect to the task (18–20). Conversely, the temporal characteristics of responses to sustained stimuli are not typically explored in detail. In fact, some have claimed that transients at stimuli onset/offset have no neuronal origin and can be explained simply in hemodynamic (21) or even nonphysiological terms: namely, step transitions in the magnetic field itself (22). However, recent studies have started to challenge this view (5, 6, 17, 23–25). Consideration of transient responses at sustained stimuli onset/offset, as well as different types of modulations (e.g., ramp, initial overshoot, and so forth) while the stimuli is present, has allowed researchers to separate responses to different stimuli type in the visual (6, 17) and auditory cortex (5, 23). Similarly, topographical dissociation between sustained and transient responses has been reported for a variety of tasks, with transient responses corresponding primarily to attention/task switching regions (24, 25). Finally, differences in transient responses between normal and clinical populations have also been reported (26). All this evidence suggests that considering a diversity of BOLD responses might help us group regions according to their relationship to the task.

To test this hypothesis, we input impulse responses computed for each subject using all 100 available runs to two different clustering algorithms (*k*-means and hierarchical clustering) that pose zero constraints on the spatial distribution of clusters. Anatomically and functionally meaningful parcellations resulted, strongly suggesting that these regions, although not showing the classic positively sustained response to the stimulus/task, are, in fact, activated.

Results

Behavioral Results. Average correct responses per run were above 95% for all subjects (Subject 1: 95.65 ± 4.06%; Subject 2: 98.12 ± 3.12%; Subject 3: 97.15 ± 3.78%), implying that subjects were compliant with the task for the entire duration of these experiments.

fMRI Activation Extent. Within-subject averaging reduced random noise but kept signal levels (nonrandom signal components) unaffected. Statistically significant signal changes (sometimes as low as less than 0.2%) time-locked with the task could be observed in almost every location of the brain (Fig. 2) when 100 runs were averaged. Response shape and magnitude varied significantly across regions. Some regions responded in a sustained positive manner for the whole duration of the task epochs (e.g., occipital, insular and left motor cortex), but others responded more prominently during task-switching periods (e.g., occipito-parietal junction). Additionally, many regions responded with negative deflections during active epochs (e.g., some parietal locations, right motor cortex). It is important to notice that even within regions that responded similarly (e.g., sustained), there are subtle differences in onset, offset, and steady-state shape. For example, although both the primary visual cortex (6 and 12 in Fig. 2) and the anterior insula (2 and 8 in Fig. 2) showed a clear sustained response for the whole duration of the active epochs, the response in the visual cortex was smoother and lacked the series of peaks and valleys clearly observable for the anterior insular region.

The extent of activations increased significantly with number of runs (N_{runs}) inputted to the analysis and relaxation of predictive BOLD response model constraints. Fig. 3 shows how activation extent rapidly increased with initial increases in number of runs ($1 < N_{runs} < 20$). For larger N_{runs}, activation extent kept increasing at a lower rate. For all subjects, significantly active voxels at $N_{runs} = 100$ represent on average over 71% of the imaged brain for the SUS model; and over 89% for the other two models at $P_{FDR} < 0.05$ (see Table S1 for results at different thresholds). Conversely, for $N_{runs} = 5$, which represents a typical number of runs per condition in fMRI experimentation, activated voxels represented ~20% of the imaged volume at $P_{FDR} < 0.05$ for the SUS analysis and between 35 and 44% for the other two analyses (Fig. 3 and Table S1).

Functional Parcellations of the Brain in Action. Voxels with similar response profiles were spatially clustered using both *k*-means and hierarchical clustering on the voxel-wise BOLD responses calculated using all 100 runs per subject. Parcellations were computed only for cortical and subcortical gray-matter voxels, excluding the

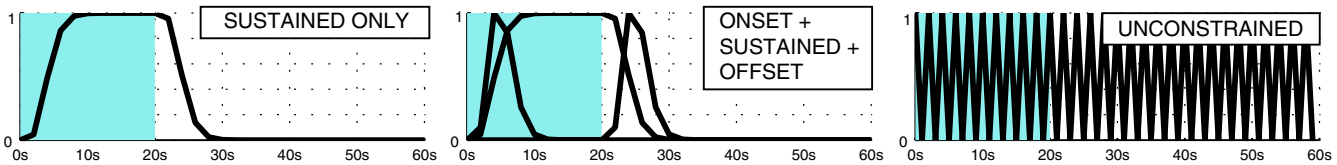

Fig. 1. Response models. Graphs span a single cycle (60 s) of the task. Active epochs (0–20 s) marked in cyan.

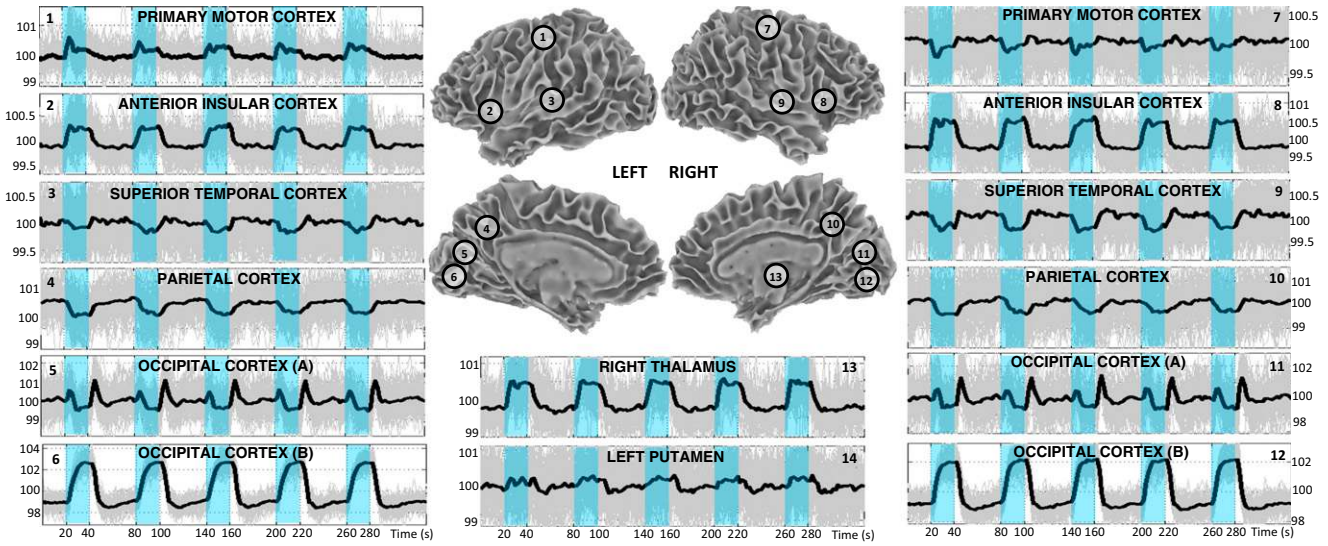

Fig. 2. Time-series for a subset of statistically significant voxels. For each voxel the 100 individual measures are plotted in gray and their average in black.

cerebellum. Parcellations were created for different clustering levels (k) ranging from $k = 2$ to $k = 70$ (see *SI Methods* for details).

Fig. 4 shows the k-means decomposition for subject 3 and $k = 20$ (Fig. S1 shows equivalent results for hierarchical clustering). In both cases, the resulting topography is symmetrical across hemispheres, anatomically meaningful, and reproducible across subjects (see Figs. S2–S4 for k-means results for all subjects, and Figs. S5–S7 for hierarchical clustering). Hemispheric symmetry is evident in the occipital cortex, superior temporal cortex, anterior insula, hippocampus, and in subcortical structures, such as the thalamus and the putamen. Moreover, clusters resemble well-known principles about the functional organization of the brain. Primary visual and primary hand motor cortices correspond to different clusters (see axial slices 21S and 49S in Fig. 4 and Fig. S1). The visual cortex is segmented into several regions both in the anterior-posterior (A-P) and medial-lateral (M-L) directions. For example, in the M-L direction V1 and V5 are segregated (see axial slice 1S in Fig. 4 and Fig. S1). In the A-P direction, V1 and higher visual processing areas closer to the parieto-occipital junction are also part of different clusters (see axial slice 9S in Fig. 4 and Fig. S1). In most cases, clusters did not appear in the form of a single contiguous agglomeration of voxels but as distributed sets of nodes. Grouping patterns go beyond hemispheric symmetry, and in some cases resemble connectivity patterns similar to those present in resting-state data. For example, CL03 in Fig. 4 (CL04 in Fig. S1) resembles a motor control network with nodes in the left primary motor hand cortex, medial supplementary motor cortex, and postero-lateral thalamus. Another example is cluster CL02 in Fig. 4 (CL05 in Fig. S1), with nodes in the bilateral infero-lateral parietal cortex, posterior cingulate, and ventro-medial frontal cortex, which resembles the default-mode network. Finally, Fig. 4*C* and Fig. S1*C* shows cluster-averaged responses. All clusters display responses time-locked with the experimental paradigm. Some clusters show positively correlated sustained responses (e.g., CL03, CL04 and CL06 in Fig. 4); others show negatively correlated sustained responses (e.g., CL01, CL02, and CL05 in Fig. 4); still others seem to respond primarily at transitions (e.g., CL09, CL11, and CL16 in Fig. 4). The average cophenetic correlation distance (CCPC) associated with the hierarchical clusters was 0.81 ($CCPC_{Subj1} = 0.85$; $CCPC_{Subj2} = 0.84$; $CCPC_{Subj3} = 0.73$), suggesting the clusters represent truly underlying structure in the data and are not artificially imposed by the clustering algorithm.

To evaluate if the clustering breaks at higher clustering orders we generated parcellations at different k levels up to a maximum

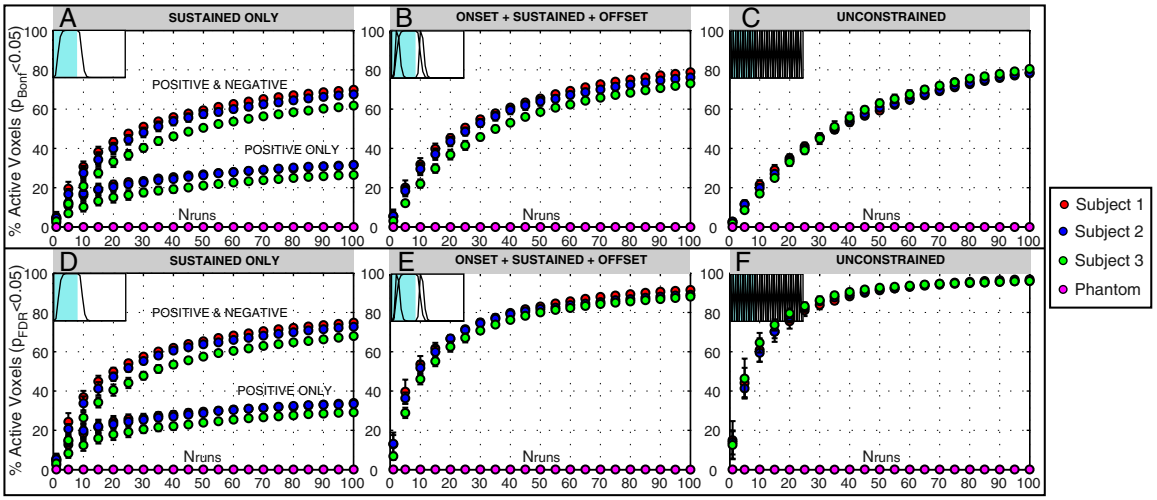

Fig. 3. Activation extent results for the three response models. *A–C* show results for $P_{Bonf} < 0.05$ and *D–F* for $P_{FDR} < 0.05$. Significantly active volume increased with number of runs in all subjects and models.

of $k = 70$ (Figs. S2–S7). Resulting functional topography at higher k levels conserved hemispherical symmetry, as well as the many functional organization principles described above.

Discussion

Using intensive within-subject averaging on data that combines blocks of visual stimulation with a number/letter discrimination task, we show that simple experimental perturbations modulate the BOLD signal in the majority of the brain. Moreover, we demonstrate that subtle interregional differences in these modulations contain sufficient information to functionally parcellate the brain at many different scales.

Noise and Predicted BOLD Response Model Significantly Affect Activation Area. The effects of averaging (8, 9) and increased versatility of response models (5, 6) on activation area have been studied separately before. Here, we evaluated their combined effects. Although additional averaging tends to increase the extent of preexisting clusters, changes in predictive BOLD response model produce novel activation sites. This finding is in agreement with previous literature. Quantitatively, noise levels have a higher effect on activation area. The maximum difference in activation area between the SUS and UNC models across all N_{runs} levels was 39% at $P_{FDR} < 0.05$. Conversely, the minimum difference between $N_{\text{runs}} = 1$ and 100 across all models was 65%.

Perhaps the most salient finding is not that activation extent increased with decreasing noise and across models, but that it reached levels close to 100% (i.e., 96%) suggesting the presence of task-based BOLD signal modulations across the whole brain. BOLD measurements are known to correlate with direct measures of neuronal activity (27); still, their relationship is not fully understood (28). Moreover, as of today there is no consensus about the neuronal interpretation of some nonconventional response shapes (e.g., deactivations, transients at stimuli onset/offset); in our case this accounts for ~68% of the reported 96% (Table S1). Nevertheless, both the 28% of voxels with conventional positive sustained responses and the remaining 68% with unconventionally shaped responses all correspond to nonrandom, task time-locked responses that reached significance at $P_{FDR} < 0.05$. Although it is perhaps premature to claim a neuronal origin for all these responses based solely on this result (see additional discussion below), it can be stated that simple, yet attention-demanding, tasks have the potential to significantly modulate on-going spontaneous BOLD fluctuations across the whole brain, including regions having no a priori relationship with the task.

Functional Organization of the Brain "in Action." Both clustering algorithms, when applied to the highest quality data, produced anatomically and functionally meaningful parcellations that were symmetrical across hemispheres, reproducible across subjects, and robust against changes in a priori selection of k or clustering methodology.

Cluster membership depended only on task-based BOLD fluctuations, because the use of 500 trials virtually eliminated all spontaneous BOLD fluctuations from the data. This finding suggests that taking into account the interregional variability of BOLD responses to sustained stimulation can provide important information about the functional organization of the brain in action, in a manner similar to that of on-going spontaneous fluctuations during rest. Nevertheless, although similarities in parcellation patterns across task-based and spontaneous fluctuations do exist, substantial differences were also present. During task performance we can observe "disconnection" of right and left motor cortex, but at rest these two areas fluctuate in synchrony (1). Furthermore, during task, right motor cortex and high order visual areas (e.g., the medio-temporal area) were assigned to the same cluster. In rest data, these regions are usually part of well-differentiated networks. This finding suggests that although the performance of a task does not completely disrupt the functional organization of the brain at rest, it does affect connectivity patterns and leads to specific rearrangements to enable performance of the task. Disruption of resting-state connectivity patterns have proven beneficial in advancing our understanding of a series of cognitive disorders, such as Autism (4) or Attention Deficit Disorder (3). We believe that disruption of normal connectivity patterns associated with specific tasks (e.g., memory-retrieval tasks or social-interaction tasks) might have similar applications. Moreover, we believe a task-based approach might be more sensitive to differences across populations because the use of tasks related to the disorder might accentuate the disrupted connectivity patterns underlying the symptoms.

Origin of Observed BOLD Signal Modulations. The use of BOLD responses other than those positively sustained as indirect markers of neuronal activity is still under debate. Both hemodynamic and neuronal mechanisms have been proposed to explain negatively sustained BOLD responses (29–31). Similarly, three models have been postulated regarding the potential origin of stimulus onset/offset responses. Two of these models are physiologically based—either hemodynamic (21) or neuronal-based (17)—and the third model is nonphysiological (22). Although our results cannot be used to categorically discard or validate any of these mechanisms, they constitute strong supporting evidence in favor of a neuronal contribution for all these categories.

One strong argument in support of a neuronal origin for non-sustained responses is spatial segregation (6, 24, 25). Our cluster results fulfill this requirement, because different responses cluster into well-delineated distinct regions. Moreover, clusters are not simply aggregations of proximal voxels, but form networks of distributed nodes with known functional homogeneity (e.g., the right primary motor and medial supplementary motor cortex are part of the same cluster). We believe these highly organized functionally based associations could not be explained solely on the basis of the vasculature tree or motion artifacts, such as those shown by Mezer et al. (32). Similarly, because the clusters of sustained negative responses cover entire, distinct brain regions, they represent more than simply neural inhibition or vasoconstriction adjacent to areas of neural excitation. Finally, if the nonphysiological model of transient responses constitutes a valid explanation, data with very low noise should be able to detect transients in most voxels. Instead, we observe some regions with only transients and others with no transients. Although it is possible for some transients to be a result of slight variations in the relationship between cerebral blood flow, blood volume, and metabolism, it is not clear how these variations would occur in some regions, but not others, and how they could account for the wide range of response shapes we observe across the brain.

Implications for Analysis/Interpretation of fMRI Data. The current observation that BOLD responses are so widely spread across the brain is unique, but it is not surprising when one considers the success of multivariate methods at decoding complex stimuli, even with regions showing no typical BOLD response to the task. Moreover, from a neuroscience perspective, this result helps narrow the gap between thousands of fMRI manuscripts showing limited activation in response to tasks and cognition theories that defend that cognition—understood as the process of "configuring the way in which sensory information becomes linked to adaptive responses and meaningful experiences"—can only result from the distributed collaboration of primary sensory, upstream and downstream unimodal, heteromodal, paralimbic, and limbic regions (33). In this context, it can be argued that noise reduction by a factor of six allowed us to switch from a regime where activity detection relates primary to sensory processing to a more sensitive regime, where activity detection includes also cognitive processes with subtler BOLD signatures.

Hypothetically, if 30% of the noise in fMRI data could be eliminated (e.g., by means of better cardiac and respiratory removal techniques), the quantity of data required to achieve noise levels equivalent to those reached here for $N_{\text{runs}} = 100$ could be cut by an order of magnitude. This result would mean that widespread task-locked activations could be detected for most tasks; therefore, this raises a fundamental question in our interpretation of fMRI data. To date, most BOLD experiments

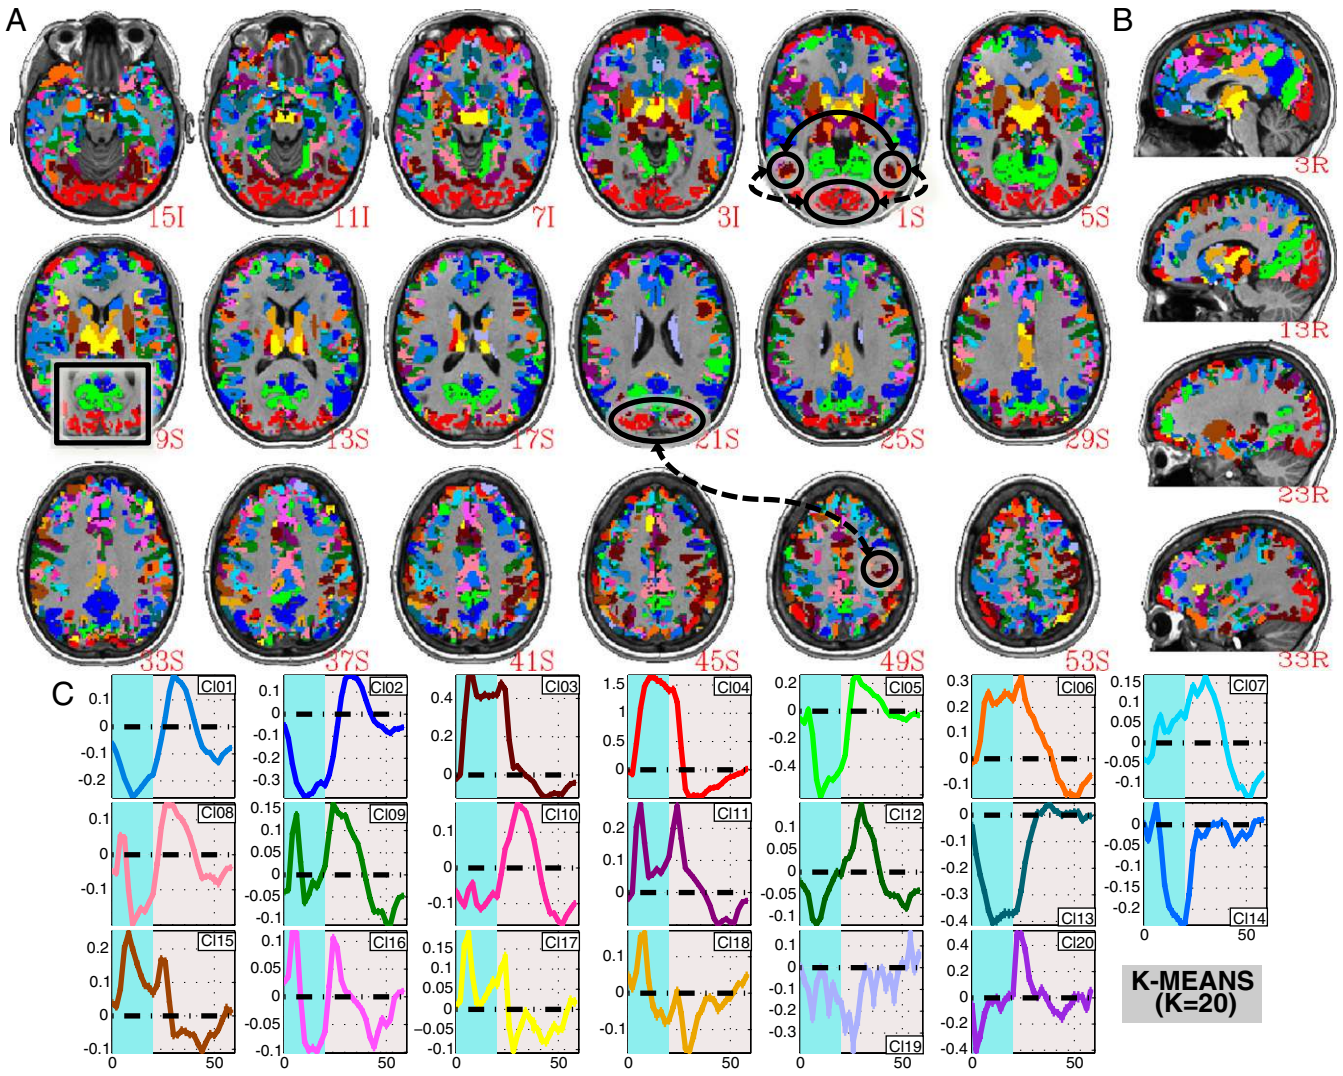

Fig. 4. Color-coded (*A*) axial and (*B*) sagittal views of *k*-means clusters ($k = 20$) for subject 3. (*C*) Color-coded hemodynamic responses for each cluster.

result in a sparse pattern of activation reflecting regions strongly responding with subtle variants on the gamma-variate hemodynamic-response function model (12). This finite number of regions facilitates the interpretation of brain function and particularly effective network modeling. What is one to make then of our current interpretation of BOLD data if this sparseness is an artifact of noise levels and constrained predictive-response models? We do not argue that whole-brain activation is critical for the processing of a certain task. However, if a task-driven BOLD response is triggered across the whole brain, how does one differentiate between BOLD responses from regions critical for handling the task, versus regions that are not? Relying on high contrast-to-noise ratio, gamma-like hemodynamic responses might not necessarily be the optimal solution. For example, language tasks generate bilateral responses reliably in fMRI (34, 35); still, a unilateral lesion can cause aphasic symptoms (36, 37). In that sense, one region is critical for the performance of the task, but the BOLD response does not indicate which.

One first step toward addressing the above-mentioned issue is to avoid the use of a simple dichotomy (e.g., active or inactive) when classifying voxels with respect to a task. If, as our data suggests, most intracranial imaged voxels should be assigned the active label under optimal experimental conditions, then such a simple dichotomy is no longer informative. Conversely, classification of voxels based on data-mining techniques, which exploit subtle differences in temporal dynamics across voxels to allow greater-order categorizations, has the potential to be more explanatory. One limitation in many of these data-mining techniques is the need to set a priori the number of output levels (e.g., clusters or components) based on hypotheses about the underlying organization of the system under study. Our data suggest, in agreement with previous resting-state studies (2, 38), that the brain "in action" shows meaningful organizational principles at many different scales, and no single number is a priori more correct than another. Nevertheless, it can be expected that some parcellations will be more informative than others. For example, although low *k* parcellations are more reliable across subjects, they do not necessarily reflect regions of homogenous response to the task. Conversely, excessively high *k* parcellations may enforce artifactual subdivisions. Algorithms that try to determine optimal parcellations by maximizing some information criteria about the data (see ref. 39 for additional details) may help address this question.

The diversity of responses observed in this work also poses interesting questions for results from univariate analyses conducted using the classic sustained-response model. In addition to underestimation of area of activation, fitting an incomplete model can be misleading because voxels with completely different shapes can produce similar fit coefficients (Fig. S8). This finding can have important implications for subtraction paradigms. Let us imagine a study in which activity for task *A* (listen to intelligible sentences)

is contrasted with task *B* (listen to unintelligible sentences) to find the neuronal correlates of intelligibility. If a given voxel were to respond in a positively sustained manner (Fig. S8*A*) to task *A*, but in a transient manner (Fig. S8*B*) to task *B*, the contrast task *A* vs. task *B* for that voxel could be nonsignificant, as Fig. S8*C* shows. By only looking to the contrast between the fits, one could conclude that such a voxel is not sensitive to speech intelligibility although the original BOLD time series clearly show that activity in this voxel highly depends on it.

Limitations of the Study. First, one limitation of this study is that the spatial resolution used here can produce partial volume effects at the edges of tissue compartments. These effects have the potential to cause an overestimation of activation extent because of the sparseness of the minimum unit of volume (e.g., in a significantly active voxel, maybe only 30% of the tissue is active; still, the entire 100% is accounted for). Higher spatial resolutions may produce smaller activation extents. Nevertheless, the fact that active tissue is distributed all over the brain—beyond areas considered as directly related to the task—and that elevated noise levels combined with generally overly simplified predictive response models produce underestimations of activation extent, remains generally valid.

A second set of limitations derives from the potential bias introduced by the selection of a specific clustering technique. To partially overcome this issue, we show results from two different methods. Although results across clustering algorithms do not exactly match on a voxel-by-voxel basis, they are highly similar. More importantly, all claims in this article are supported by results from either approach. In other words, no claims that could be supported only with results from a single method were made.

Finally, a third limitation of this study is the impossibility to unquestionably claim that the effects reported here—namely the presence of task time-locked BOLD responses in the majority of the brain in response to a simple stimulation + letter/number discrimination task—are truly neuronal in origin. Although anatomical correspondence and functional relevance of reported clusters represent strong arguments in favor of a significant neuronal contribution, conclusive evidence to support such a claim would require simultaneous recording of BOLD signals and neuronal events from the entire brain, with combined simultaneous EEG and fMRI recording being a plausible option.

Conclusions

We have demonstrated that the sparseness of activations in fMRI maps can result from elevated noise levels or overly strict predictive BOLD response models. When noise is sufficiently low and the response model versatile enough, activity can be detected with BOLD fMRI in the majority of the brain. Finally, we have demonstrated that subtle interregional differences in BOLD response shapes contain sufficient information to produce functional parcellations of the brain "in action."

Methods

Three subjects were scanned on a General Electric 3T MRI scanner. All subjects underwent 100 functional runs, which consisted of five blocks of stimulation (20 s: flickering checkerboard at 8 Hz + letter/number discrimination task) and 40 s of rest. For the letter/number discrimination task, subjects responded using a response box with their right hand. Data were analyzed with AFNI (prepreprocessing, and statistical analysis), and MATLAB (clustering). Motion and physiological noise fluctuations were removed during preprocessing. Active voxels are defined as those where the model accounts for a significant amount of variability in the data (*F*-stat) at $P_{FDR} < 0.05$ or $P_{Bonf} < 0.05$.

1. Biswal B, Yetkin FZ, Haughton VM, Hyde JS (1995) Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn Reson Med* 34: 537–541.
2. Smith SM, et al. (2009) Correspondence of the brain's functional architecture during activation and rest. *Proc Natl Acad Sci USA* 106:13040–13045.
3. Castellanos FX, et al. (2008) Cingulate-precuneus interactions: A new locus of dysfunction in adult attention-deficit/hyperactivity disorder. *Biol Psychiatry* 63:332–337.
4. Cherkassky VL, Kana RK, Keller TA, Just MA (2006) Functional connectivity in a baseline resting-state network in autism. *Neuroreport* 17:1687–1690.
5. Harms MP, Melcher JR (2003) Detection and quantification of a wide range of fMRI temporal responses using a physiologically-motivated basis set. *Hum Brain Mapp* 20: 168–183.
6. Uludağ K (2008) Transient and sustained BOLD responses to sustained visual stimulation. *Magn Reson Imaging* 26:863–869.
7. Lieberman MD, Cunningham WA (2009) Type I and Type II error concerns in fMRI research: Re-balancing the scale. *Soc Cogn Affect Neurosci* 4:423–428.
8. Saad ZS, Ropella KM, DeYoe EA, Bandettini PA (2003) The spatial extent of the BOLD response. *Neuroimage* 19:132–144.
9. Huettel SA, McCarthy G (2001) The effects of single-trial averaging upon the spatial extent of fMRI activation. *Neuroreport* 12:2411–2416.
10. Calhoun VD, Stevens MC, Pearlson GD, Kiehl KA (2004) fMRI analysis with the general linear model: Removal of latency-induced amplitude bias by incorporation of hemodynamic derivative terms. *Neuroimage* 22:252–257.
11. Mechelli A, Henson RN, Price CJ, Friston KJ (2003) Comparing event-related and epoch analysis in blocked design fMRI. *Neuroimage* 18:806–810.
12. Cohen MS (1997) Parametric analysis of fMRI data using linear systems methods. *Neuroimage* 6:93–103.
13. Lindquist MA, Wager TD (2007) Validity and power in hemodynamic response modeling: A comparison study and a new approach. *Hum Brain Mapp* 28:764–784.
14. Henson RN, Price CJ, Rugg MD, Turner R, Friston KJ (2002) Detecting latency differences in event-related BOLD responses: Application to words versus nonwords and initial versus repeated face presentations. *Neuroimage* 15:83–97.
15. Handwerker DA, Ollinger JM, D'Esposito M (2004) Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. *Neuroimage* 21:1639–1651.
16. Aguirre GK, Zarahn E, D'esposito M (1998) The variability of human, BOLD hemodynamic responses. *Neuroimage* 8:360–369.
17. Tyler CW, Kontsevich LL, Ferree TC (2008) Independent components in stimulus-related BOLD signals and estimation of the underlying neural responses. *Brain Res* 1229: 72–89.
18. Smolders A, et al. (2007) Dissecting cognitive stages with time-resolved fMRI data: A comparison of fuzzy clustering and independent component analysis. *Magn Reson Imaging* 25:860–868.
19. Formisano E, Goebel R (2003) Tracking cognitive processes with functional MRI mental chronometry. *Curr Opin Neurobiol* 13:174–181.
20. d'Avossa G, Shulman GL, Corbetta M (2003) Identification of cerebral networks by classification of the shape of BOLD responses. *J Neurophysiol* 90:360–371.
21. Obata T, et al. (2004) Discrepancies between BOLD and flow dynamics in primary and supplementary motor areas: Application of the balloon model to the interpretation of BOLD transients. *Neuroimage* 21:144–153.
22. Renvall V, Hari R (2009) Transients may occur in functional magnetic resonance imaging without physiological basis. *Proc Natl Acad Sci USA* 106:20510–20514.
23. Giraud AL, et al. (2000) Representation of the temporal envelope of sounds in the human brain. *J Neurophysiol* 84:1588–1598.
24. Fox MD, Snyder AZ, Barch DM, Gusnard DA, Raichle ME (2005) Transient BOLD responses at block transitions. *Neuroimage* 28:956–966.
25. Konishi S, Donaldson DI, Buckner RL (2001) Transient activation during block transition. *Neuroimage* 13:364–374.
26. Fox MD, Snyder AZ, McAvoy MP, Barch DM, Raichle ME (2005) The BOLD onset transient: Identification of novel functional differences in schizophrenia. *Neuroimage* 25:771–782.
27. Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A (2001) Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412:150–157.
28. Logothetis NK (2008) What we can do and what we cannot do with fMRI. *Nature* 453: 869–878.
29. Harel N, Lee SP, Nagaoka T, Kim DS, Kim SG (2002) Origin of negative blood oxygenation level-dependent fMRI signals. *J Cereb Blood Flow Metab* 22:908–917.
30. Shmuel A, Augath M, Oeltermann A, Logothetis NK (2006) Negative functional MRI response correlates with decreases in neuronal activity in monkey visual area V1. *Nat Neurosci* 9:569–577.
31. Devor A, et al. (2007) Suppressed neuronal activity and concurrent arteriolar vasoconstriction may explain negative blood oxygenation level-dependent signal. *J Neurosci* 27:4452–4459.
32. Mezer A, Yovel Y, Pasternak O, Gorfine T, Assaf Y (2009) Cluster analysis of resting-state fMRI time series. *Neuroimage* 45:1117–1125.
33. Mesulam MM (1998) From sensation to cognition. *Brain* 121:1013–1052.
34. Gonzalez-Castillo J, Talavage TM (2011) Reproducibility of fMRI activations associated with auditory sentence comprehension. *Neuroimage* 54:2138–2155.
35. Binder JR, et al. (1997) Human brain language areas identified by functional magnetic resonance imaging. *J Neurosci* 17:353–362.
36. Bates E, et al. (2001) Differential effects of unilateral lesions on language production in children and adults. *Brain Lang* 79:223–265.
37. Broca PP (1861) Remarques sur la siege de la faculte du langage articule, suivies d'une observation d'aphemie (perte de la parole) [Remarks on the seat of the faculty of articulated language, following an observation of aphemia (loss of speech)]. *Bulletin de la Societe d'Antrophologie* 2:330–357. French.
38. Damoiseaux JS, et al. (2006) Consistent resting-state networks across healthy subjects. *Proc Natl Acad Sci USA* 103:13848–13853.
39. Goutte C, Hansen LK, Liptrot MG, Rostrup E (2001) Feature-space clustering for fMRI meta-analysis. *Hum Brain Mapp* 13:165–183.