NeuroImage 49 (2010) 1949–1952

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

Comments and Controversies

Spatial smoothing hurts localization but not information: Pitfalls for brain mappers

Yukiyasu Kamitani *, Yasuhito Sawahata

ATR Computational Neuroscience Laboratories, 2-2-2 Hikaridai, Keihanna Science City, Kyoto 619-0288, Japan

ARTICLE INFO

Article history:
Received 3 April 2009
Revised 12 June 2009
Accepted 17 June 2009
Available online 25 June 2009

ABSTRACT

Op de Beeck (Op de Beeck, H., 2009. Against hyperacuity in brain reading: Spatial smoothing does not hurt multivariate fMRI analyses? Neuroimage) challenges the possibility of extracting information from subvoxel representations via random biases associated with voxel sampling, the hypothesis proposed by Kamitani and Tong (Kamitani, Y., Tong, F., 2005. Decoding the visual and subjective contents of the human brain. Nat. Neurosci. 8, 679–685). Here, we show that his results provide no evidence against the possibility, being consistent with both of the subvoxel and supravoxel representation models. Classification of spatially smoothed fMRI data is not an effective means to probe into information sources for multivoxel decoding, since smoothing does not hurt the information contents of multivoxel patterns. We point out the danger of interpreting multivoxel decoding results based on intuitions guided by the conventional brain mapping paradigm.

In Kamitani and Tong (2005), we proposed the notion of 'ensemble feature selectivity', the selectivity achieved through optimally combining fMRI voxels allowing for prediction or decoding of stimulus features such as orientation and motion direction. Such selectivity can arise from voxels each of which is only weakly selective due to the lack of the spatial resolution of the standard fMRI to resolve the putative feature representations. We suggested that information could be extracted from random biases associated with voxel sampling of subvoxel neural and vascular structures, which are known to have substantial irregularities.

The issue of the current debate raised by Op de Beeck is whether random biases of subvoxel structures can account for the observed decoding results, or some broader supravoxel representations should underlie the decoding. We acknowledge that this is an essential issue regarding the foundation of multivoxel decoding analysis, and requires further investigation. However, the argument put forward by Op de Beeck in an attempt to challenge the possibility of subvoxel sources seems to be based on an incorrect belief about the mathematical nature of spatial smoothing, and on an artifact of the simplified simulation. Here, we critically examine the reasoning of Op de Beeck's argument and the underlying assumptions, and show that his results are orthogonal to the issue of debate, being consistent with both of the subvoxel and supravoxel information sources. We specifically aim to remove common misconceptions regarding information represented in multivoxel patterns, while leaving more substantive arguments on the actual information sources for multivoxel decoding to other authors.

* Corresponding author. Fax: +81 744 95 1259.
E-mail address: kmtn@atr.jp (Y. Kamitani).

1053-8119/$ – see front matter © 2009 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2009.06.040

No loss of information by spatial smoothing

Op de Beeck found that spatial smoothing using a Gaussian kernel does not hurt decoding performance, and took this observation as evidence against the theory of subvoxel sources. His argument apparently assumes that spatial smoothing degrades information represented in fine-scale patterns. Is this assumption true?

Let an original voxel pattern $\mathbf{x}$ ($N \times 1$ vector; N, number of voxels), and a smoothed voxel pattern $\mathbf{x}'$ ($N \times 1$). Smoothing can be expressed by

$$\mathbf{x}' = \mathbf{K}\mathbf{x}, \tag{1}$$

where $\mathbf{K}$ represents the smoothing kernel: each row represents the smoothing weights for each element of $\mathbf{x}'$. If smoothing does not involve downsampling, $\mathbf{K}$ is a square matrix with each row having a spatially shifted elements. As each row is linearly independent of the other rows, $\mathbf{K}$ is full-ranked. Hence, $\mathbf{K}$ is invertible. The original voxel pattern $\mathbf{x}$ can be recovered from the smoothed pattern by

$$\mathbf{x} = \mathbf{K}^{-1}\mathbf{x}'. \tag{2}$$

As shown above, smoothing, or more generally convolution, is an invertible transformation, unless there is downsampling or a complete cutoff of high frequency components (note that a Gaussian kernel does not involve complete cutoffs). The inverse transformation, or deconvolution, is known to be not robust in the presence of noise added after smoothing. But the present case does not involve such noise (except for some numerical errors). Thus, a smoothed voxel pattern can be transformed back to the original voxel pattern without loss of information. An example of the complete recovery of a smoothed image is shown in Fig. 1a. This demonstrates that even though fine-scale features are smeared by smoothing, they can be completely recovered by the inverse transformation.

Next, let us consider a binary linear classification problem. Supposed that we obtain an optimal linear discriminant function,

$$f(\mathbf{x}) = \mathbf{w}^{\mathrm{T}}\mathbf{x}, \tag{3}$$

where $\mathbf{w}$ is a vector representing voxel weights (bias terms included in $\mathbf{w}$ and $\mathbf{x}$; T, transpose). The classification boundary between two classes is defined by the hyperplane $f(\mathbf{x})=0$. Test samples are classified according whether $f(\mathbf{x})>0$ (class 1) or $f(\mathbf{x})<0$ (class 2).

We now consider classification in the smoothed voxel space. By substituting $\mathbf{K}^{-1}\ \mathbf{x}'$ for $\mathbf{x}$ in $f(\mathbf{x})$, we have a linear discriminant function for a smoothed voxel pattern $\mathbf{x}'$

$$g(\mathbf{x}') = \mathbf{w}^{\mathrm{T}}\mathbf{K}^{-1}\mathbf{x}' = \mathbf{v}\mathbf{x}' \tag{4}$$

where $\mathbf{v}$ is the weight vector defined by $\mathbf{w}^{\mathrm{T}}\ \mathbf{K}^{-1}$. Note that this discriminant function $g(\mathbf{x}')$ always gives the same classification result as $f(\mathbf{x})$, that is, if $f(\mathbf{x})>0$, then $g(\mathbf{x}')>0$, and vice versa. Therefore, the original and the smoothed data are equivalent, in the sense that one can obtain identical classification results using an optimal discriminant function. More generally, it can also be shown that the likelihood ratio between the distributions for the two classes remains the same after spatial smoothing at any corresponding points in the multivoxel space, indicating that the degree of overlap between the distributions is not affected by smoothing.

A practical question is whether it is better to estimate $\mathbf{w}$ of $f(\mathbf{x})$ in the original voxel space, or $\mathbf{v}$ of $g(\mathbf{x}')$ in the smoothed voxel space, given a limited number of training data. This is a matter of how well the data fit the nature of the mathematical model and the estimation algorithm. The accurate classification obtained with different levels of smoothing (Fig. 3 in Op de Beeck (2009)) seems to suggest that SVM is good at estimating the parameters both in the original and the smoothed space.

In actual fMRI data analysis, a subset of voxels are often selected as a region of interest (ROI) after the smoothing of the whole brain volume. Within the ROI, this smoothing may not be invertible, because signals outside the ROI are involved. Hence, for the comparison of the original and smoothed patterns, the contamination of non-ROI voxels should be carefully examined.

Smoothing with an appropriate scale may be important for localizing relevant representations, as Op de Beeck notes. However, as discussed above, smoothing does not hurt information, and an optimal classifier can perform the same level of decoding regardless of the degree of smoothing. Thus, the decoding results of smoothed data need to be carefully interpreted when the spatial scale of information sources is discussed.

Voxels are not independent in motion-corrected fMRI data

Besides the high decoding performance maintained after smoothing, Op de Beeck's argument is based also on the size of correlation between fMRI activity patterns for the same stimulus. He found that the correlation size increased with the degree of smoothing, and that this trend was consistent only with the simulation result from a large-scale (supravoxel) representation, but not with the result from a small-scale (subvoxel) representation, in which smoothing did not affect the correlation size. Op de Beeck took this observation as evidence for the existence of some large-scale representation underlying the fMRI data.

First, it should be noted that the analysis of the correlation size does not tell us how distinct the voxel patterns for different stimuli

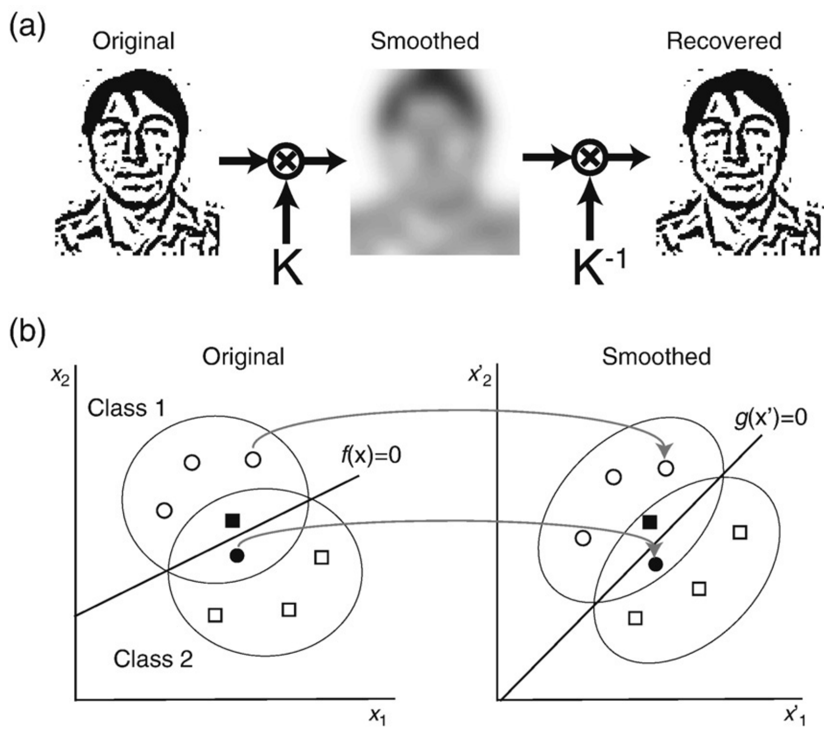

Fig. 1. No loss of information by spatial smoothing. (a) Complete recovery of a smoothed image. An image on our lab's web site portraying one of the authors was used as the input (120 × 109 pixels, 8 bit gray scale). Sharp edges were emphasized by applying the 'nancyKOseki filter' (http://www.hirax.net/misc/nancyKOseki/) to the source photo. Smoothing was performed by multiplying the image vector (13080 [= 120 × 109] × 1) by the kernel matrix ($\mathbf{K}$; 13080 × 13080), in which each row represented 2D Gaussian weights (standard deviation, 50 pixels, isotropic). The smoothed image was then transformed by the inverse matrix $\mathbf{K}^{-1}$. (b) Schematic representation of voxel patterns before and after smoothing. Each symbol represents an fMRI activity pattern (circles, class 1; squares, class 2) on a two-voxel space (x_1–x_2, original; x'_1–x'_2, smoothed). Arrows indicate the mapping from $\mathbf{x}$ to $\mathbf{x}'$ caused by the smoothing. $f(\mathbf{x})$ and $g(\mathbf{x}')$ denote linear discriminant functions performing exactly the same classification (open symbols, correctly classified; filled symbols, misclassified).

are. When the correlation size within a class increases, the correlation size across classes can also increase. As we have seen, the information about stimulus classes is preserved after smoothing, regardless of the correlation size.

Then, why did smoothing have different effects on the simulation results for the small- and the large-scale representations? A critical difference is in the (in)dependency of neighboring voxels sampled from these representations. The signals of neighboring voxels from the small-scale representation are independent, being determined by the biases of randomly arranged ‘cells’ within each voxel region. In contrast, the signals of neighboring voxels from the large-scale representation are dependent as they tend to share the same preference. In the presence of signal dependency between neighbors, smoothing may well help attenuate noise, resulting in a large correlation size between voxel patterns. Note that in this simulation, it is assumed that voxels are sampled from exactly the same regions in every volume scan.

In actual fMRI measurement, however, it is difficult to achieve such exact voxel samplings, because of the subject's head motions and other sources of image shifts. It is customary to perform a motion-correction procedure (Ashburner and Friston, 2004), as is done in Op de Beeck's analysis of real fMRI data. But this procedure disrupts the independence of voxel sampling. As illustrated in Fig. 2a, a residual head motion smaller than the voxel size is often corrected by linear interpolation using the neighboring voxels. A spline-based interpolation method (default option in SPM) involves an even broader range of voxels to determine the intensity of a voxel. Thus, the simulation assuming independent voxel sampling is not realistic.

We performed a slightly more realistic simulation, in which we repeated the procedure of Op de Beeck's simulation except that voxel samplings were randomly shifted every trial, and motion correction was applied to create the data for analysis. Results show qualitatively similar trends between the small and the large-scale representations in both correlation size and classification performance (Fig. 2b, motion correction by linear interpolation). Notably, the correlation size increased with the degree of smoothing even for the small-scale representation, in contrast to the result of Op de Beeck's simulation. The difference between the small- and large-scale representations became even smaller when a spline-based motion correction was applied. Thus, both of the small- and the large-scale simulation results are consistent with the real fMRI results, providing no ground for ruling out the possibility of extracting information form a small-scale representation.

In motion-corrected fMRI data, each voxel represents a broader cortical region than a single voxel volume, and thus the signals of neighboring voxels are not independent even for a small-scale representation. The accurate classification obtained with the motion-corrected simulation data (Fig. 2b; see also Fig. 6 in Kamitani and Tong (2005)) suggests that each voxel can preserve reliable information for decoding by detecting statistical biases in a region broader than a single voxel volume.

Does size matter?

It may worth mentioning another prevailing assumption, which also sounds intuitive but is not necessarily true. It is often argued that

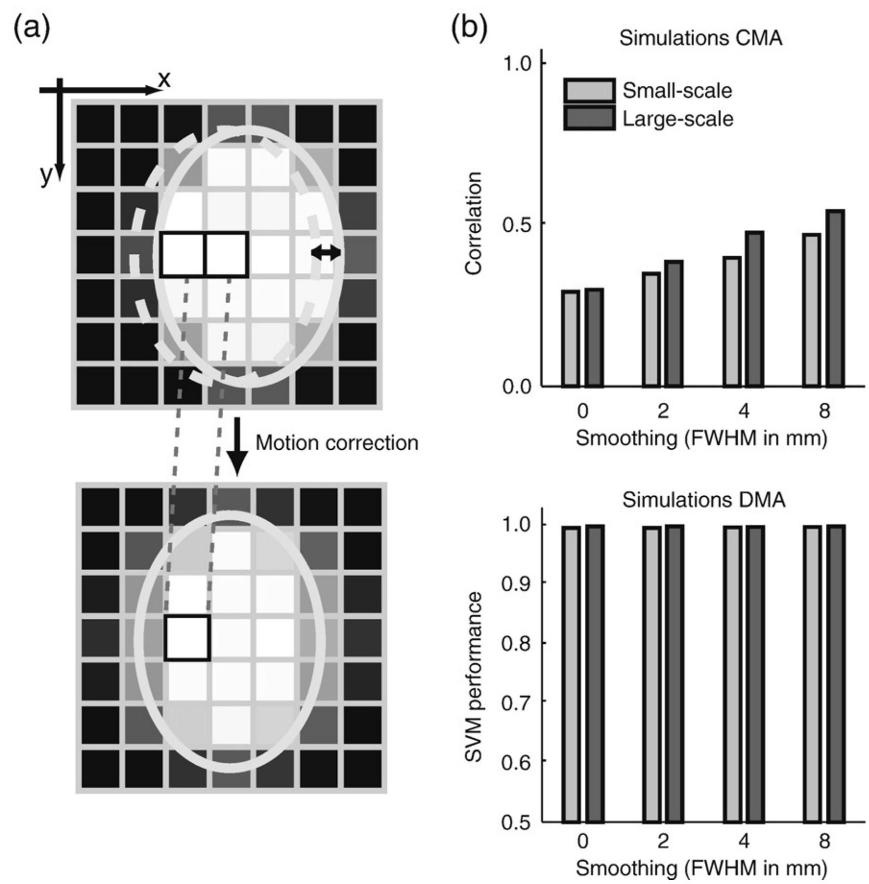

Fig. 2. Voxels are not independent in motion-corrected fMRI data. (a) Schematic view of motion correction. The upper and lower panels present an original and a motion-corrected fMRI images, respectively. The solid oval in the upper panel shows the actual position of the head. The dashed oval indicates the target position to which all images should be aligned. The head motion is corrected by linear interpolation using neighboring voxels. Resulting voxel intensities are no longer independent of the adjacent ones. (b) Simulation results with motion correction. A random subvoxel head motion was introduced in each data sample, and then it was corrected by the standard linear interpolation method. Other procedures were the same as those in Op de Beeck's simulations.

a larger voxel should be less informative about stimuli represented in subvoxel units, and therefore that if the decoding using large voxels (often created by downsampling) maintain a high level of accuracy, it indicates the presence of some large-scale representation. Downsampling, unlike spatial smoothing, is a non-invertible transformation and thus can degrade information. However, if one considers a single voxel or the same number of voxels, a large voxel (or N large voxels) can be as informative as a small voxel (or N small voxels). Here, we only hint at this point by providing a simple example (more rigorous discussion will be given in our separate paper in preparation).

Consider a cortical representation consisting of two types of cells preferring to either stimulus 1 or stimulus 2. These cells respond to the stimuli with the same activity levels (+0.5 for preferred stimulus; −0.5 for non-preferred stimulus) plus independent noise ($\sim N(0, \sigma^2)$). If each voxel contains M cells (motion correction not considered here), and the cell's preference is randomly assigned with a probability of 0.5, the number of cells preferring to stimulus 1 (M_1) follows a binomial distribution, $M_1 \sim B(M, 0.5)$, having a mean $\mathrm{E}(M_1) = M/2$ $(= M \times 0.5)$ and a variance $\mathrm{Var}(M_1) = M/4$ $(= M \times 0.5 \times (1 - 0.5)$; the same distribution for 'stimulus 2 cells', M_2; $M_1 + M_2 = M$).

In the analysis of subvoxel representations, the bias $M_1 - M_2$ $(= 2\,M_1 - M)$ is assumed to serve as the signal in each voxel. Using $\mathrm{Var}(M_1) = M/4$, the variability of the signal, the square root of $\mathrm{Var}(M_1 - M_2) = \mathrm{Var}(2\,M_1 - M)$ becomes $M^{1/2}$. On the other hand, as the voxel size (M) increases, the noise in each voxel, that is, the sum of independent noise from M cells, increases in proportion to $M^{1/2}$, the same rate as the signal variability. Therefore, at least in this simple model, a large voxel can be as informative as a small voxel on average. This example provides a cautionary note against the intuition, 'the larger the less informative about small representations'. Careful inspection will be necessary in interpreting decoding results obtained with different voxel sizes.

Conclusion

We have critically examined the assumptions behind Op de Beeck's reasoning, and shown that his results provide no piece of evidence against 'hyperacuity in brain reading'. Such assumptions may seem intuitive from the mapping or localization point of view, but they need to be carefully scrutinized when the information represented in multivoxel patterns is discussed. Although it may be possible that spatial smoothing could be effectively used to reveal information sources underlying fMRI decoding, it would require more sophisticated and quantitative modeling and study design.

In this comment, we have discussed voxel sampling as if fMRI signals directly derive from neural responses. But this by no means presents a realistic picture. fMRI is an indirect measure of neural signals mediated by hemodynamic responses, and large vessels are known to substantially contribute to the signals (Turner, 2002). The vasculature may play the role of another 'biased sampler' of neural representations, constituting a nested vascular-voxel sampling process. It is important to note that the principle presented in the previous section ('Does size matter?') may apply to the vascular sampling, too. As the collective response from a broad cortical region can be as informative as that from a small region, it may be possible that large vessels that drain blood from a broad cortical region can carry reliable information about stimulus features represented in small neural structures.

Acknowledgments

The authors thank O. Yamashita for helpful comments. This work was supported in part by Nissan Science Foundation and SRPBS, MEXT.

References

Ashburner, J., Friston, K.J., 2004. Rigid body registration. In: Frackowiak, R.S.J. (Ed.), Human Brain Function. Academic Press, San Diego, pp. 635–654.
Kamitani, Y., Tong, F., 2005. Decoding the visual and subjective contents of the human brain. Nat. Neurosci. 8, 679–685.
Op de Beeck, H., 2009. Against hyperacuity in brain reading: spatial smoothing does not hurt multivariate fMRI analyses? Neuroimage.
Turner, R., 2002. How much cortex can a vein drain? Downstream dilution of activation-related cerebral blood oxygenation changes. Neuroimage 16, 1062–1067.