

Mining Local Connectivity Patterns in fMRI Data

Kristian Loewe¹, Marcus Grueschow², and Christian Borgelt³

Abstract A core task in the analysis of functional magnetic resonance imaging (fMRI) data is to detect groups of voxels that exhibit synchronous activity while the subject is performing a certain task. Synchronous activity is typically interpreted as functional connectivity between brain regions. We compare classical approaches like statistical parametric mapping (SPM) and some new approaches that are loosely based on frequent pattern mining principles, but restricted to the local neighborhood of a voxel. In particular, we examine how a soft notion of activity (rather than a binary one) can be modeled and exploited in the analysis process. In addition, we explore a fault-tolerant notion of synchronous activity of groups of voxels in both the binary and the soft/fuzzy activity setting. We apply the methods to fMRI data from a visual stimulus experiment to demonstrate their usefulness.

1 Introduction

The localization and analysis of brain activity is a major objective in cognitive neuroscience. Functional magnetic resonance imaging (fMRI) provides an indirect, but non-invasive means to measure brain activity in vivo. Essentially, time series of three-dimensional (3D) brain-images are acquired, in which each volumetric pixel (or *voxel* for short) represents a cuboid of tissue. Inferences about brain activity rest on the following principle: neuronal activity entails the consumption of oxygen and thus the supply of the

¹ Department of Knowledge and Language Processing, University of Magdeburg, D-39106 Magdeburg, Germany, kristian.loewe@gmx.net

² Laboratory for Social and Neural Systems Research, Department of Economics, University of Zürich, CH-8006 Zürich, Switzerland, marcus.grueschow@econ.uzh.ch

³ European Centre for Soft Computing, c/ Gonzalo Gutiérrez Quirós s/n, E-33600 Mieres (Asturias), Spain, christian.borgelt@softcomputing.es

relevant area with oxygenated blood. The different magnetic properties of oxygenated blood in comparison to deoxygenated blood result in observable signal changes in the time series of the relevant voxels, which are exploited as an indirect indicator of neuronal activity. This is known as the blood oxygen level dependent (BOLD) effect [9]. For an excellent review of neurovascular coupling and its effect on the BOLD signal see [8].

Typical task-related fMRI experiments are designed and conducted on the grounds of certain hypotheses about brain functions which are subsequently tested using regression-based statistics. To this end, most often general linear models (GLM) based on canonical hemodynamic response functions (cHRF) are fitted to each individual voxel time series in order to obtain statistical maps highlighting brain activity related to experimental conditions. In that regard, several *a priori* assumptions are widely accepted by the neuroimaging community. For example, to facilitate comparison between voxels, GLMs are generated and fitted using the same cHRF for all time series, even though hemodynamic responses differ widely across the brain [1, 5]. As a consequence, such an analysis is limited to the testing of *a priori* hypotheses and frequently makes use of *a priori* assumptions, which—in case they are not met—may constrain the significance of the obtained results.

In contrast to this, data-driven approaches might reveal unexpected patterns that in turn could give rise to new hypotheses, while at the same time *a priori* assumptions are avoided as far as possible. For example, recent studies made use of graph-theory in order to derive network characteristics of the brain from interregional functional connectivity matrices, where functional connectivity means the temporal dependence between brain regions [4]. In both, task-related and resting-state settings (where in a resting state no task and no explicit external stimulus is presented) studies indicated that the brain network is organized in a highly clustered way [3, 13]. Recently, this was exploited to compute locally restricted correlations (based on spatial proximity) in order to rapidly identify potential hub regions in the brain [11].

Building in a similar fashion on the strongly clustered brain organization, we present a new, noise-robust, and purely data-driven method targeting local functional connectivity patterns. The proposed approach is applicable to any type of fMRI data (task-related, resting-state etc.) and allows for time-efficient and model-free generation of meaningful brain maps (without making *a priori* assumptions or presuming hypotheses to test).

2 Notions of Activity

fMRI data are series of periodically acquired 3D intensity images. We denote one such series by $\mathbf{i} = (i_1, i_2, \dots, i_T)$, where T is the number of points in time at which the intensity images i_k , $k \in \{1, \dots, T\}$, are recorded. The individual images are organized in a regular 3D voxel grid of size $X \times Y \times Z$.

In order to simplify the processing, the voxel coordinates (x, y, z) can be mapped (in an essentially arbitrary, but fixed fashion) to a linear index v with $1 \leq v \leq V = X \cdot Y \cdot Z$. In this way a data set can be represented by a data matrix $\mathbf{S}^{V \times T} = (s_{v,t})$. By $\mathbf{s}_v = s_{v,*} = (s_{v,1}, s_{v,2}, \dots, s_{v,T})$, that is, the v -th row of \mathbf{S} , we denote the time series of voxel v .

As the data are arbitrarily scaled, a meaningful and comparable notion of activity (magnitude) arises only from a relative interpretation. In the following, we derive a binary and a soft notion of voxel activity by considering at any given time the deviation from its temporal average intensity. The binary notion can be seen as a limiting case of the soft notion.

Binary Notion of Activity. We use a simple binary discretization in order to assign to each value in a time series one of the two qualitative states *active* and *inactive*. Formally, the dichotomized time series d_v of a voxel v is given by $\mathbf{d}_v = (d(s_{v,1}), d(s_{v,2}), \dots, d(s_{v,T})) \in \{0, 1\}^T$ induced by the function $d(s_{v,t}) = H(s_{v,t} - \tilde{s}_v)$, $t \in \{1, \dots, T\}$, where \tilde{s}_v denotes the median of the values in \mathbf{s}_v and H is the Heaviside step function, defined as $H(x) = 0$ if $x < 0$ and $H(x) = 1$ otherwise. In other words, a voxel is regarded as active at a point in time if the corresponding signal intensity value amounts at least to the median of the respective time series. The median was chosen over the mean because it is less sensitive to outliers.

Note that this very simple scheme is naturally open to many points of criticism. For example, it enforces that a voxel is active half of the time, which is clearly debatable. However, it already leads to useful results and thus we defer finding better discretization schemes to future work.

Note also that the concise binary time series representation \mathbf{d}_v can be exploited in order to speed up subsequent analysis through a highly efficient implementation using bit vectors. However, this advantage comes at the expense of the inevitable loss of information due to the discretization.

Soft Notion of Activity. The above discretization implies an extreme sharpening of the signal: whereas the actual signal rises gradually over time, the discretization enforces a sharp instantaneous signal change once the median is exceeded. Effectively, the signal is transformed into a square-wave signal, thus increasing the contrast at the transition sites.

By replacing the Heaviside step function with a sigmoid function (for instance, a logistic function), we introduce a soft notion of activity, which enables a parameterized sharpening of the signal (thus also limiting the information loss). Formally, we transform the time series according to the linear scaling and logistic activation function

$$f_{\text{act}}^{(\beta)}(s_{v,t}) = \left(1 + \exp \left(-\frac{s_{v,t} - \tilde{s}_v}{\beta(Q_{0.95}(\mathbf{s}_v) - Q_{0.05}(\mathbf{s}_v))} \right) \right)^{-1},$$

where $Q_{0.05}(\mathbf{s}_v)$ and $Q_{0.95}(\mathbf{s}_v)$ denote the 5% and the 95% quantile, respectively, of the time series \mathbf{s}_v . Their difference can be seen as an estimate of the




Fig. 1 Application of logistic activation functions differing in their slope β to a time series \mathbf{s} . For $\beta \rightarrow \frac{1}{\epsilon}$ this is equivalent to the discretization approach.

range of intensity values, which is more robust than simply using min and max and thus is in line with our choice of the median over the mean.

The strength of the sharpening effect is governed by the slope parameter β . Obviously the binary discretization is obtained as a limiting case of this scheme for $\beta \rightarrow \frac{1}{\epsilon}$. An illustration is shown in Figure 1.

Note that we use the normalization by an estimate of the range of values in order to keep the meaning of the slope parameter independent of the range of values of the time series. Of course, this is also open to criticism, as it removes all information related to the amount of signal change, which may contain valuable information. However, as with the choice of the median as the transition point between inactive and active, we leave further improvements of the activation scheme for future work.

3 Local Connectivity Measures

Recent voxelwise functional connectivity analyses showed a highly clustered organization of the brain in both task-related as well as resting-state settings [3, 13]. Aiming to characterize local connectivity patterns by quantifying the local cohesion strength, we propose new local connectivity measures (LCM) operating on the time series of the enclosing $3 \times 3 \times 3$ cuboid of each voxel. In this scheme each center voxel serves as an identifier of its enclosing cuboid, allowing for the data to be traversed in a sliding 3D window fashion. For this purpose, we denote by

$$N_{26}(v) = \{w \mid v \neq w \wedge \max\{|x_v - x_w|, |y_v - y_w|, |z_v - z_w|\} \leq 1\}$$

the 26 neighbors of a voxel $v = (x_v, y_v, z_v)$ (formed by the 6 voxels sharing a face, the 12 sharing an edge and the 8 sharing a vertex with it). The whole

cuboid of a voxel v is then denoted by $C_{27}(v) = \{v \cup N_{26}(v)\}$ and we only consider voxels with a complete cuboid neighborhood.

A natural approach would be to calculate the average of Pearson's correlation coefficient of all $\frac{27 \cdot 26}{2} = 351$ voxel pairs within an enclosing cuboid. However, each individual correlation—and thus also their mean—might be considerably prone to noise. Especially if there are some voxels that do not participate in the joint activity and thus have low correlations with the other voxels, an existing co-activity pattern may not be discernible.

Therefore, instead of combining pairwise correlations by averaging, we try to obtain a more robust measure by integrating the values of all 27 voxels at each point in time. To be more precise, our idea is to find for each voxel's enclosing cuboid $C_{27}(v)$ those points in time that exhibit synchronous activity of at least a certain (user-specified) number of voxels. Formally, we have

$$\text{LCM}_\alpha^B(v) = \frac{1}{T} \sum_{t=1}^T H\left(\left(\sum_{w \in C_{27}(v)} d_{w,t}\right) - \alpha\right),$$

where H is the Heaviside step function. “LCM” stands for “local connectivity measure” and the upper index B indicates that it is based on the binary time series \mathbf{d}_v . The measure is normalized w.r.t. T , the length of the time series, in order to facilitate comparison between data sets of different length. The parameter $\alpha \in \{1, \dots, 27\}$ captures the fault-tolerant aspect of this measure: Given a cuboid and a point in time, α active voxels suffice for the corresponding addend to become equal to one (and thus to contribute positively to the co-activity measure).

For functionally independent adjacent voxels we expect to see around 13–14 active voxels at each point in time, because with our discretization scheme (active above and inactive below the median) each voxel is active at half of the points in time and therefore about half of the voxels in a cuboid should be active on average. As a consequence, α should be chosen greater than 14. However, what choice of α is best depends on the level of noise present in the data and the desired contrast between functionally connected and functionally independent 27-cuboids.

Clearly, the number of active voxels can be expected to be significantly higher than 14 at points in time actually showing co-activity. As this co-activity “uses up” some of the active states of the participating voxels, the remaining points in time must possess a lower average number of co-active voxels. In addition, it is plausible to assume that functionally connected voxels also exhibit co-inactivity, that is, possess points in time at which only a relatively low number of voxels are active. This can be exploited to enhance the contrast of the measure by defining

$$\text{LCMd}_\alpha^B(v) = \text{LCM}_\alpha^B(v) + (1 - \text{LCM}_{28-\alpha}^B(v)).$$

Since $\text{LCM}_\alpha^B(v)$ is the higher, the more co-activity the voxels in a cuboid show, while $(1 - \text{LCM}_{28-\alpha}^B(v))$ is the higher, the more co-inactivity they show, this measure can be expected to be more sensitive.

Using our soft notions of activity, we now define soft analogues of the above measures. In order to handle the activity degrees, we rely on the following reasoning: in a perfect situation, in which all voxels are active at a point in time ($\alpha = 27$), the terms summed over can also be seen as conjunctions of the activity values (active — 1, inactive — 0). In a soft setting this conjunction could be expressed, using a standard fuzzification of conjunctions, by a minimum. The fault-tolerant aspect can then be incorporated by replacing the minimum with a quantile, thus allowing a few activations to be low. This leads to the following measure:

$$\text{LCM}_{\alpha,\beta}^S(v) = \frac{1}{T} \sum_{t=1}^T Q_{1-\frac{\alpha-1/2}{27}} \left(\left[f_{\text{act}}^{(\beta)}(s_{w,t}) \mid w \in C_{27}(v) \right] \right),$$

where $Q_p([x_1, \dots, x_k])$ denotes the p -quantile of the data set $[x_1, \dots, x_{27}]$ (which we do not write as a set in order to allow for multiple voxels having the same activation) such that $1 - \frac{\alpha-1/2}{27}$ selects the α -smallest value. “LCM” again stands for “local connectivity measure” and the upper index S indicates that it is based on a soft notion of activity.

Arguing in the same way as for the binary measure, we can increase the contrast and thus the sensitivity for detecting co-activity by defining

$$\text{LCMd}_{\alpha,\beta}^S(v) = \text{LCM}_{\alpha,\beta}^S(v) + (1 - \text{LCM}_{28-\alpha,\beta}^S(v)).$$

4 Data and Preprocessing

We applied the proposed methods to both artificial data and real fMRI recordings from a task-related experiment.

Artificial Data. In order to analyze the characteristics of the new measures, we generated synthetic data sets of co-active and independent voxels. One sample of a data set consisted of 27 voxel time series corresponding to one $3 \times 3 \times 3$ cuboid of voxels. The co-active samples were created using zero vectors (of length 300) into which blocks of ones were inserted at random locations. The vectors were then convolved with the cHRF included in the software package SPM8¹ for Matlab². Finally, white Gaussian noise (WGN) of given signal-to-noise ratio (SNR) was added. In this way, three data sets of co-active samples were created, corresponding to SNRs of +10, 0 and

¹ SPM8. Wellcome Trust Centre for Neuroimaging, London, UK.
Available at <http://www.fil.ion.ucl.ac.uk/spm/>.

² MATLAB®. The MathWorks Inc., Natick, Massachusetts, USA.

−10 decibel (dB). In addition, we created one data set containing samples of independent voxels using WGN time series.

Real Data. The usefulness of the proposed methods as applied to real data was assessed using task-related fMRI data, acquired using a 7 Tesla MR-scanner (Siemens, Erlangen, Germany) in the context of a study of the visual pathway. Subjects were instructed to focus on a central fixation point, while being exposed to alternating left and right visual hemifield stimulation of different luminance contrasts. Meanwhile, functional data were acquired in volumes of $192 \times 192 \times 27$ voxels at isotropic resolution of 1.1mm edge length using a time resolution of 2s. Details can be found in [12].

The analysis of fMRI data is susceptible to manifold artifacts arising from both physiological and hardware-related sources. It was therefore essential to account for them prior to the actual analysis. Using an online image-reconstruction procedure, all data were motion- and distortion-corrected based on a reference measurement of the local point spread function [14]. As interpolation causes local correlations not originally present in the data, we refrained from spatial smoothing and normalization to a standard brain.

Frequencies below 0.01Hz were removed from the individual voxel time series accounting for low frequency signal intensity drifts caused e.g. by scanner instabilities [10] and physiological artifacts. Non-brain voxels were excluded from further analysis by defining a brain mask using a thresholding procedure based on the means of the voxel time series. For comparative purposes in further analyses the remaining brain voxels were also partitioned into gray matter and non-gray matter voxels by thresholding of a gray matter probability map generated using SPM8 segmentation routines.

GLM Analysis. For comparisons, a conventional GLM analysis of the task-related fMRI data was carried out using SPM8. We applied a statistical model containing boxcar waveforms convolved with a cHRF, representing the left and right visual hemifield stimulation, respectively. Multiple linear regression was then used to generate parameter estimates for each regressor at every voxel. Visual field biased regions in each subject were identified using a contrast of contralateral greater than ipsilateral visual stimulation resulting in a statistical parametric map of t -statistics (SPM_t).

5 Results for Test Data

Three variants of LCM and LCMd were calculated for the four test data sets (see Section 4): LCM_α^B , $LCMd_\alpha^B$, $LCM_{\alpha,0.05}^S$, $LCMd_{\alpha,0.05}^S$, $LCM_{\alpha,0.1}^S$, and $LCMd_{\alpha,0.1}^S$ (Figure 2). As anticipated, the LCM is higher for the co-active voxels than for the noise voxels for $\alpha > 14$ while the opposite is true for $\alpha \leq 14$. LCMd exploits both complementary contrasts and thus exhibits an increased sensitivity compared to LCM.




Fig. 2 The generated test data (Section 4) were subject to LCM_{α}^B (top left), $LCMd_{\alpha}^B$ (top right), $LCM_{\alpha,0.05}^S$ (mid left), $LCMd_{\alpha,0.05}^S$ (mid right), $LCM_{\alpha,0.1}^S$ (bottom left), and $LCMd_{\alpha,0.1}^S$ (bottom right). Sample means corresponding to the three data sets consisting of co-active 27-cuboids (circles) and to the WGN 27-cuboids (asterisks) were plotted against α . For the former, the darkness of the gray decreases with the SNR used when adding WGN to the time series. The absolute differences between the results corresponding to the WGN data and those corresponding to the three co-active data sets were plotted adopting the respective gray levels (triangles). The same holds for the horizontal dashed lines representing the mean of the average correlation coefficient (of the 351 pairwise correlations per sample) distribution of the respective co-active data set.




Fig. 3 Distribution of local connectivity measure LCMd for three signal-to-noise ratios.

The average correlation coefficient yields higher contrast between co-active and WGN samples (as estimated by the difference of their respective results) if the strength of the noise is low, that is, for higher SNRs. The opposite is true for lower SNRs: here LCM/LCMd outperforms the average correlation coefficient in the course of decreasing SNR. Unavoidably, however, the higher the strength of the additive WGN, the more the LCM/LCMd of the co-active samples resemble those of the WGN data.

As illustrated by Figure 3, the variance (and range) of LCMd—and therefore its sensitivity—is lowest for the most extreme values of α . Accordingly, also the difference between the co-active and the WGN voxels is minimal for $\alpha = 14$ and $\alpha = 27$ (Figure 2, right column). As explained in Section 3, the best choice of α depends on the level of noise present in the data and the desired contrast between functionally connected and functionally independent 27-cuboids. While a smaller α provides higher noise robustness (fault tolerance), a too small choice will impair the contrast between co-active voxels and WGN voxels, as the expected value under the assumption of independent voxels comes closer. Then again, some fault tolerance needs to be ensured, as due to noise a large α will result in a low range and similar LCMd for both co-active and noise voxels, all the same.

The attainable noise robustness of LCM seems to increase with the sharpening effect, that is, with decreasing slope β . The smaller β is, the farther each value in a time series gets shifted towards minimum or maximum, that is, towards 0 or 1. Thus, with β decreasing, LCM values tend to increase for $\alpha > 14$ and to decrease for $\alpha \leq 14$. In other words, the soft approach seems to keep more noise than actual information. However, this behavior may be due to our scheme of transforming the time series (using a median and quantile normalization) and further investigations are needed in order to clarify this.

6 Results for Real Data

The model-driven approach using a GLM yielded robust activation in cortical and subcortical visual regions mostly confined to gray matter. As expected, highest activity differences were found in the left and right calcarine sulcus respectively, the location of retinotopically organized primary visual cortex (V1). Additional regions representing the left and right visual field, respectively, could be localized such as the secondary and mid-level visual regions V2 adjacent to the dorsal and ventral part of V1 as well as area MT+, located bilaterally in the temporal parts of occipital cortex. Subcortical regions such as the lateral geniculate nucleus (LGN) as well as the superior colliculus showed statistically significant differences between their left and right visual field representation, albeit much lower than V1 (see Section 4).

We now set out to address the question whether it is possible to generate meaningful brain maps without any *a priori* assumptions with respect to experimental design or hemodynamic properties across the brain. The LCM analysis generally yielded higher values for gray matter regions than for non-gray matter regions. As for SPM_t , the highest values were found in left and right V1 and adjacent visual regions. Visual inspection of the results indicated that high SPM_t values are most often accompanied by high LCM values (3rd and 4th row of Figure 4). The two-dimensional histogram of SPM_t and LCM confirmed this observation (5th row of Figure 4). Conversely, many voxels were exhibiting no considerable activation associated with the visual experiment, while at the same time showing a high LCM. For both hemispheres, two regions of interest (ROI) of 100 voxels each were defined in the center of GLM activation, i.e., in V1, and in a white matter (WM) region in the temporal lobe, where no coherent—let alone visually driven—activity was to be expected. Highly active voxels in both V1 ROIs, identified with GLM and indicated by high SPM_t deviations from zero, were also identified by LCM, while both approaches identified the WM regions as non-responsive. Voxels identified as highly active by SPM_t also show high LCM values, suggesting local connectivity increases as the cortex is active (5th row of Figure 4).

7 Discussion, Conclusions and Future Work

We presented noise-robust and data-driven measures that characterize local functional connectivity patterns in fMRI data. Specifically, the proposed LCMs are designed in order to capture the proportion of synchronous activity (LCM, $\alpha > 14$), synchronous inactivity ($1 - \text{LCM}$, $\alpha < 14$) or both (LCMd) as exhibited by adjacent voxels during a fixed period of time.

Using fMRI data from a study of the visual path, we compared stimulus-related activity as detected by conventional regression-based GLM analysis with local functional connectivity as estimated by LCMd. While increased

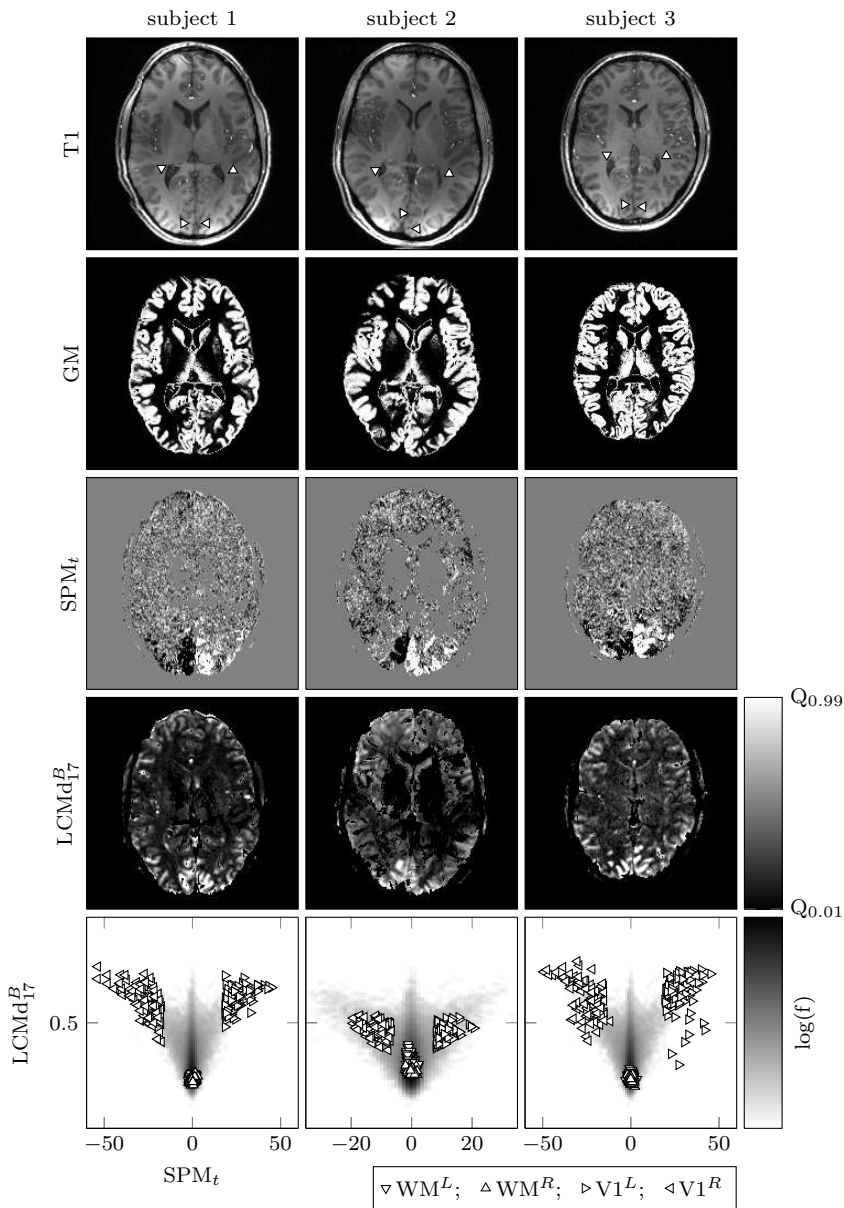


Fig. 4 GLM and LCM results. Each column corresponds to one subject from the visual stimulation experiment (Section 4). For each panel in row 1-4 the underlying data were clipped and scaled according to its respective 0.01- and 0.99-quantile before an axial slice was mapped to a gray color scale. 1st row: T1-weighted anatomical image. 2nd row: gray matter (GM) probability map. 3rd row: GLM results (SPM_t). Difference map between activity caused by ipsi- and contra-lateral visual field stimulation. 4th row: $LCMd_{17}^B$ map. 5th row: SPM_t vs. $LCMd_{17}^B$. Background: 2D histogram of all brain voxels whose GM probability exceeded 0.5. Four ROIs of 100 voxels each are shown on top.

stimulus-related activity was most often accompanied by increased local functional connectivity, we also detected functionally connected clusters that exhibited no considerable task-related activity. These clusters may have been affected by locally coherent noise patterns or they may have been engaged in neuronal activity unrelated (or at least not linearly related) to the visual stimuli. In the latter case, further inspection of these clusters might give rise to new hypotheses, subsequently testable in a conventional fashion.

Beyond the initial proof of concept, the proposed approach may be utilized in future neuroscientific research as well as possible therapeutical implementations. As the method is entirely data-driven, it is applicable to *any* fMRI data (task- or stimulus-induced, resting-state, etc.). As such, LCM-based analyses may be especially suited to the analysis of resting-state fMRI data, as in this case no experimental task or stimulus onsets exist on which regressors for GLM and fitting of cHRFs could be based.

With a properly adjusted implementation, the LCM-versions based on dichotomous time series (that is, with a binary notion of activity) allow for time-efficient analysis of very large and very many data sets. This aspect might be exploitable for real-time fMRI (rtfMRI). The rtfMRI methodology aims at efficiently analyzing neuroimaging data in an online fashion (that is, concurrently with the data acquisition by the scanner), the results of which may govern the adaptation of experimental stimulation and the interaction with the subject. The feasibility of online analysis of complex emotional and cognitive states has recently been shown [7], while the future aim of such methods lies in therapeutic neurofeedback-based training after traumatic brain injury, cognitive stress or neurological pathology and will potentially culminate in brain machine interfaces [6]. In this application domain, changes in local functional connectivity could serve as an indicator of changing activity patterns, as suggested by the comparison of GLM/SPM_t and LCM results.

In addition, LCM might serve as a filter in order to constrain the brain voxels to be analyzed further based on the functional images only or in addition to a T1-based gray matter segmentation. In favor of this idea it can be said that the LCM maps and the GM probability maps seemed to be highly conform (which is not surprising, though, since no neuronal activity is to be expected in non-GM areas). In fact, an initial and general reduction to informative parts of fMRI data before the actual analysis (whether assumption free or not) may constitute an interesting field of potential applications.

Future work includes finding a better way of mapping the intensity signal as picked up by the scanner to an activation degree, since the shortcomings of our current mapping do not allow us to fully exploit the advantages of a soft approach, which inherently is better suited to maintain all relevant information. Secondly, we are working on perturbation schemes to generate surrogate data that can be used to derive *p*-values for the detected local connectivity. Finally, we are in the process of extending our approach to a time-efficient analysis of spatially *unconstrained* connectivity, which is made possible by bit-vector representations of a binary notion of activity.

Acknowledgements The work presented in this paper was supported by Short-Term Scientific Mission (STSM) grant 9059 (Kristian Loewe) of COST Action IC0702.

References

1. Aguirre GK, Zarahn E, D’Esposito M (1998) The Variability of Human, BOLD Hemodynamic Responses. *Neuroimage* 8(4):360–369. Elsevier, Amsterdam, Netherlands
2. Cordes D, Haughton VM, Arfanakis K, Carew JD, Turski PA, Moritz CH, Quigley MA, Meyerand ME (2001) Frequencies Contributing to Functional Connectivity in the Cerebral Cortex in “Resting-state” Data. *American Journal of Neuroradiology* 22(7):1326–1333. American Society of Neuroradiology, Oak Brook, IL, USA
3. Eguiluz VM, Chialvo DR, Cecchi GA, Baliki M, Apkarian AV (2005) Scale-free Brain Functional Networks. *Physical Review Letters* 94:18102. American Physical Society, College Park, MD, USA
4. Friston KJ, Frith CD, Liddle PF, Frackowiak RS (1993) Functional Connectivity: The Principal-component Analysis of Large (PET) Data Sets. *Journal of Cerebral Blood Flow and Metabolism* 13(1):5–14. Nature Publishing Group, London, United Kingdom
5. 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(4):1639–1651. Elsevier, Amsterdam, Netherlands
6. Hollmann M, Mönch T, Mulla-Osman S, Tempelmann C, Stadler J, Bernarding J (2008) A New Concept of a Unified Parameter Management, Experiment Control, and Data Analysis in fMRI: Application to Real-time fMRI at 3T and 7T. *Journal of Neuroscience Methods* 175(1):154–162. Elsevier, Amsterdam, Netherlands
7. Hollmann M, Rieger JW, Baecke S, Lützkendorf R, Müller C, Adolf D, Bernarding J (2011) Predicting Decisions in Human Social Interactions using Real-time fMRI and Pattern Classification. *PLoS ONE*, 6(10):e25304. Public Library of Science, San Francisco, CA, USA
8. Logothetis NK, Wandell BA (2004) Interpreting the BOLD Signal. *Annual Review of Physiology* 66:735–769 Annual Reviews, Palo Alto, CA, USA
9. Ogawa S, Lee TM, Nayak AS, Glynn P (1990) Oxygenation-sensitive Contrast in Magnetic Resonance Image of a Rodent Brain at High Magnetic Fields. *Magnetic Resonance in Medicine* 14(1):68–78. J. Wiley & Sons, Chichester, United Kingdom
10. Smith AM, Lewis BK, Ruttimann UE, Ye FQ, Sinnwell TM, Yang Y, Duyn JH, Frank JA (1999) Investigation of Low Frequency Drift in fMRI signal. *Neuroimage* 9(5):526–533. Elsevier, Amsterdam, Netherlands
11. Tomasi D, Volkow ND (2010) Functional Connectivity Density Mapping. *Proceedings of the National Academy of Sciences of the USA* 107(21):9885. National Academy of Sciences, Washington, DC, USA
12. Tschukalin A (2011) Noninvasive Lokalisation von magno- und parvozellulären Anteilen des humanen CGL mittels Hochfeld-MRT. Bachelor Thesis. Dept. of Computer Science, Otto-von-Guericke Universität Magdeburg, Germany
13. Van den Heuvel MP, Stam CJ, Boersma M, Hulshoff Pol HE (2008) Small-world and Scale-free Organization of Voxel-based Resting-state Functional Connectivity in the Human Brain. *Neuroimage* 43(3):528–539. Elsevier, Amsterdam, Netherlands
14. Zaitsev M, Hennig J, Speck O (2004) Point Spread Function Mapping with Parallel Imaging Techniques and High Acceleration Factors: Fast, Robust, and Flexible Method for Echo-planar Imaging Distortion Correction. *Magnetic Resonance in Medicine* 52(5):1156–1166. J. Wiley & Sons, Chichester, United Kingdom