Non-negative partial least squares for meta-analytic parcellation: A functional atlas for the human brain

Finn Årup Nielsen[1,2,3], Daniela Balslev[4], Lars Kai Hansen[1,2]

[1]Lundbeck Foundation Center for Integrated Molecular Brain Imaging, Copenhagen, Denmark; [2]Informatics and Mathematical Modelling, Technical University of Denmark, Lyngby, Denmark; [3]Neurobiology Research Unit, Copenhagen University Hospital Rigshospitalet, Copenhagen, Denmark; [4]Danish Research Centre for Magnetic Resonance, Copenhagen University Hospital Hvidovre, Hvidovre, Denmark.

September 19, 2006

Abstract

Algorithms can automatically analyze human brain mapping studies represented in a neuroinformatics database. We describe one such meta-analytic method that relies on a combination of text mining, functional volumes modeling, partial least squares and non-negative matrix factorization (NMF): NMF decomposes the product between a bag-of-words matrix, constructed from abstract words, and a voxelization matrix constructed by kernel density modeling of the stereotaxic coordinates in the scientific papers contained in a database. The components found allows us to construct a functional atlas where voxels and words get loaded on components interpretable as brain functions. When applied on the Brede Database with 186 papers we find components such as memory, emotion, pain and audition. Furthermore, we present a cluster validation procedure based on permutation and cluster matching that quantifies the variability of the functional atlas.

1 Introduction

One of the aims in human brain mapping is to establish a functional atlas of the entire human brain, making it possible to label each brain area with one or more specific functions. Kleist (1934) constructed such an atlas based on correlating the location of brain lesions with behavioral examinations, where many of the areas corresponded to those of Brodmann (1994), and in the modern era several atlases show the macaque visual areas (Van Essen, 2003). For more anatomy-oriented atlases the systematic collection of brain regional features allows, e.g., Kötter et al. (2001) to perform multivariate analyses of multimodal receptor binding and connectivity in macaque, and Zilles et al. (2004) and Schleicher et al. (1999) to characterize the human cerebral cortex based on multimodal receptor binding profile and cytoarchitectonic profile, respectively.

Modern human brain mapping produces a wealth of information linking brain and behavioral variables. Here we describe a system for automated construction of a functional atlas based on a neuroinformatics database.

We work with data from the Brede Database (Nielsen, 2003). As inspired from the BrainMap database (Fox et al., 1994) it contains 3-dimensional stereotaxic Talairach coordinates (Talairach and Tournoux, 1988) representing focal brain lesions or change in brain activity to the mental process under investigation as reported in published neuroimaging papers. In the Brede Database we have previously used multivariate analysis in the form of singular value decomposition, independent component analysis, K-means clustering and non-negative matrix factorization to extract consensus patterns across sets of articles represented as voxelized volumes (Nielsen and Hansen, 2004; Nielsen, 2003). Though the analyses produce maps in Talairach space, which in some cases can have an interpretation as a specific brain function, they do not directly allow us to functionally label voxels automatically. In a previous meta-analysis we sought the largest elements in the joint density between abstract words and Talairach coordinates for functional labeling of Talairach coordinates (Nielsen et al., 2004b). In that work we showed that non-negative matrix factorization provides a means for exploiting the covariance between words and voxels and in our present work we pursue and extend this so we can functionally label all grey matter areas of the human brain, either with individual words for

a a's aberrant aberrations abilities ability ablated ablations able abnormal abnormalities abnormality abolished about above ... you you'd you'll you're you've young younger your yours yourself yourselves z zero zone zones ... amygdala amygdaloid angular anterior area basal bilateral brain brainstem calcarine callosomarginalis caudal caudate ... temporo temporoparietal thalamic thalamus uncus ventral ventrolateral ventromedial ventroposterior vermis vi viib

Table 1: Excerpt of stop word list.

brain functions or with a "component" that loads on words as well as voxels.

2 Method

Our procedure combines four different methods for the parcellation of the human brain into functional areas: Text modeling, functional volumes modeling, non-negative matrix factorization and partial least squares. In each case we take the simplest approach. For our method we rely on the implementations in the Brede Toolbox (Nielsen and Hansen, 2000a).

2.1 Text modeling

The abstracts of papers in the Brede Database form the base for text modeling and we use the so-called "vector space model" (Salton et al., 1975), where a bag-of-words matrix $\mathbf{X}(N \times P)$ represents P different words in N different abstracts. The absolute frequency (the count) of the p'th word in the n'th abstract c_{np} enters the (n, p) matrix element as the square root $x_{np} = \sqrt{c_{np}}$ (Cutting et al., 1992). We choose this normalization to control the variance of each element: Regarding each count as generated from an individual Poisson distribution, the count itself can act as a fairly good estimate of the variance $\hat{\sigma}^2_{np} = c_{np}$, and with $x_{np} = c_{np}/\hat{\sigma}_{np}$ each matrix element gets standardized to approximately unit variance.

The most common words of English ("the", "of", etc.) dominate the matrix. A *stop word* list, setup in one of our previous studies (Nielsen et al., 2005), eliminates these common words as well as words for brain anatomy and a large number of other words that appear often in scientific abstracts, see Table 1 for a few examples. Words for brain function dominate the matrix after application of this list. We also eliminate words that only occur in a single abstract.

2.2 Functional volumes modeling

For the modeling of the Talairach coordinates — "functional volumes modeling" (Fox et al., 1997) — we construct a matrix $\mathbf{Y}(N \times Q)$ from voxelization of all the Talairach coordinates in each of the N papers. Our method uses kernel density estimation on the Talairach coordinates (Chein et al., 2002; Nielsen and Hansen, 2002b; Turkeltaub et al., 2002; Wager et al., 2003; Nielsen and Hansen, 2000b) convolving each 3-dimensional coordinate $\mathbf{v}_l$ in the n'th paper with a 3-dimensional isotropic Gaussian kernel in Talairach space $\mathbf{v}$

$$p(\mathbf{v}|n) = \sum_e^{E_n} \sum_l^{L_{n,e}} \frac{(2\pi\sigma^2)^{-3/2}}{\sqrt{E_n L_{n,e}}} \exp\left[\frac{(\mathbf{v}-\mathbf{v}_l)^2}{-2\sigma^2}\right], \tag{1}$$

with E_n as the number of experiments (e.g., individual contrast images) in the n'th paper and $L_{n,e}$ as the number of coordinates in each experiment. We get a pseudo-probability density $p(\mathbf{v}|n)$ since the normalization uses the square root of the numbers and thus the density does not integrate to one. The normalization tries to weight equally coordinates regardless of whether the individual coordinate comes from a paper with few or many other coordinates — weakly justified from a voting system (Penrose, 1946; Nielsen, 2005). We set the kernel width to $\sigma = 10$mm and then sample the pseudo-probability density on an $8 \times 8 \times 8$ millimeter grid obtaining a volume for each paper. Labeled voxels in the "AAL" atlas of Tzourio-Mazoyer et al. (2002) act as mask for grey matter regions, resulting in Q number of voxels. With the voxels within the mask vectorized and the row vectors for N papers stacked we get the voxelization matrix $\mathbf{Y}$.

Slight differences exist between templates used for spatial normalization, and the anatomical differences might make it difficult or impossible to match multiple brains. Thus two locations with the same coordinate do not necessarily appear in the same (anatomical) brain area if they come from two different studies. For the Brede Database the program for data entry corrects for the difference between the (original) Talairach and the Montreal Neurological Institute (MNI) spaces by applications of "Brett's transformation" (Brett, 1999), though this does not fully correct the discrepancy (Chau and McIntosh, 2005; Lancaster et al., 2006).

Some voxel-based meta-analyses choose to filter their studies, including only studies with a field of view for the entire human brain (Fox et al., 2005). The Brede Database does presently not represent the imaging or reporting field of view for the major part of the papers. It means that we cannot automate a correction for this, and results may become biased.

Instead of a voxel-based analysis we could match each location to a brain area specified, e.g., by its lobar anatomy or Brodmann area, building a (paper × brain area)-matrix (Indefrey and Levelt, 2000; Lloyd, 2000). We view this as a more complex procedure, since one would need to specify the precise anatomical delineation and construct a matching algorithm that assigns a location to an area. One would either rely on "author-labels" or "atlas-labels" (Laird et al., 2005). Neuroimaging authors applies very little standardized neuroanatomical terminology, so the match becomes difficult for author-labels (Laird et al., 2005; Nielsen, 2006). For atlas-labeling the area of an atlas label should overlap with the area of the individual study, but, e.g., what an author labels cerebellum may not fall in the cerebellum of the atlas (Nielsen and Hansen, 2002b), so one would expect some slight mislabeling.

2.3 Partial least squares

The text modeling and the functional volumes modeling give two data sets represented in matrices $\mathbf{X}$ and $\mathbf{Y}$, and our further modeling correlates these two matrices. There exists a number of multivariate analysis methods for this such as canonical correlation analysis (Mardia et al., 1979), inter-battery factor analysis (Tucker, 1958), partial least squares by singular value decomposition (McIntosh et al., 1996; McIntosh and Lobaugh, 2004), and canonical ridge analysis (Vinod, 1976; Mardia et al., 1979; Nielsen et al., 1998), all which submit to a formulation as a generalized eigenvalue decomposition (Borga et al., 1992). We will take the simplest form of these methods and use partial least squares. This first constructs a product matrix

$$\mathbf{Z} = \mathbf{X}^\mathsf{T}\mathbf{Y}. \tag{2}$$

In one type of partial least squares technique a singular value decomposition extracts interesting factors from the product matrix. However, we instead use non-negative matrix factorization on the product matrix and we term this *non-negative partial least squares* (nPLS).

2.4 Non-negative matrix factorization

Both of matrices from text modeling and functional volumes modeling are non-negative, i.e., all elements are either zero or positive. Any matrix product between two non-negative matrices becomes non-negative. One finds a number of algorithms and models specifically targeting non-negative matrices, e.g., the latent class model (Goodman, 1974), non-negative CANDECOMP (Carroll et al., 1989), factor analysis with non-negative transformation (Shen and Israël, 1989), positive matrix factorization (Paatero, 1997), non-negative matrix factorization (Lee and Seung, 1999, 2001), the "aspect model" (Saul and Pereira, 1997; Hofmann, 1999), mixture of multinomials (Rigouste et al., 2005; Nigam et al., 2000), replicator dynamics (Neumann et al., 2005), non-negative PARAFAC (Bro and Jong, 1997) and different variations of non-negative independent component analysis (Højen-Sørensen et al., 2002; Plumbley, 2003), where some of these share similar parameter update forms (Gaussier and Goutte, 2005).

Here we will not explore the full range of non-negative algorithms, but confine us to the "Euclidean" version of non-negative matrix factorization (NMF), which stands as simple, reasonable fast, row/column symmetric in its decomposition, and with no hyperparameters, except for the number of components. The product of two non-negative matrices $\mathbf{W}(P \times K) \geq 0$ and $\mathbf{H}(K \times Q) \geq 0$ acts as a low-rank approximative factorization of the product matrix $\mathbf{Z} = \mathbf{X}^\mathsf{T}\mathbf{Y}$

$$\mathbf{WH} = \mathbf{Z} + \mathbf{U}, \tag{3}$$

with $\mathbf{U}$ as a residual matrix. We use the iterative updates of Lee and Seung (2001):

$$\begin{aligned} \mathbf{H}_{kp} &\leftarrow \mathbf{H}_{kp} \frac{\left(\mathbf{W}^\mathsf{T}\mathbf{Z}\right)_{kp}}{\left(\mathbf{W}^\mathsf{T}\mathbf{WH}\right)_{kp}} \qquad (4) \\ \mathbf{W}_{qk} &\leftarrow \mathbf{W}_{qk} \frac{\left(\mathbf{Z}\mathbf{H}^\mathsf{T}\right)_{qk}}{\left(\mathbf{WH}\mathbf{H}^\mathsf{T}\right)_{qk}}. \qquad (5) \end{aligned}$$

We let the algorithm run for 50000 iterations or until no considerable change between iterations appears. We also run the factorization multiple times with different initializations and pick the one with the lowest Frobenius norm of the residual matrix. Our hardware exhibited slow performance during underflow multiplications with denormalized numbers (Goldberg, 1991), and during iterations we sat elements in $\mathbf{W}$ and $\mathbf{H}$ to zero if they where below 10^{20} times the underflow threshold. Furthermore, we add a small value to the nominators of the update formula 4 and 5 to avoid divisions by zero (Shahnaz et al., 2006).

This standard NMF fails to be unique since scaling factors can be moved between the columns of $\mathbf{W}$ and the rows of $\mathbf{H}$. We compute vectorial 2-norms and distribute the scaling equally between the two matrices. The remaining variation may be captured by our cluster validation procedure described later.

As the only parameter the number of components K needs to be set. We presently set this to $K \approx \sqrt{N/2}$ — related to the rule of thumb of Mardia et al. (1979, p. 365) — rather than try to optimize this.

For exclusive assignment of each word and each voxel to a component K we apply winner-take-all functions to $\mathbf{W}$ and $\mathbf{H}$

$$\tilde{w}_{pk} = \begin{cases} w_{pk} & \text{if } \forall_{k'} : w_{pk} \geq w_{pk'} \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

$$\tilde{h}_{kq} = \begin{cases} h_{kq} & \text{if } \forall_{k'} : h_{kq} \geq h_{k'q} \\ 0 & \text{otherwise.} \end{cases} \tag{7}$$

This will result in a two-way clustering of both words and voxels, i.e., a hard parcellation of both words and voxels.

2.5 Cluster validation

The clustering obtained with the nPLS varies depending on the data, the scaling of the matrices and the initialization of the NMF algorithm. A number of measures for quantifying the variability of (any type of) clusterings has been devised, see, e.g., (Meilă, 2002; Law and Jain, 2003). These measures typically quantify the cluster stability based on the labels from two clusterings. One type of measures require that the individual clusters should be mapped (or "matched") between the clusterings: The first cluster in the first clustering corresponds not necessarily to the first cluster in the second clustering (Strehl and Ghosh, 2002). One possible way of matching, the greedy algorithm, seeks out the largest element in the confusion matrix, say a_{ij}, and matches the i'th cluster from the first clustering with the j'th cluster from the second clustering, then deletes the i'th row and the j'th column from the confusion matrix and performs the operation again repeatedly on the reduced matrix (Meilă and Heckerman, 2001; Meilă, 2002). Another algorithm, the so-called Hungarian method (Kuhn, 1955; Roth et al., 2002; Lange et al., 2003; Tichavský and Koldovský, 2004), may make a better pairing at the expense of larger time complexity. We implemented a variation of this algorithm (Munkres, 1957; Pilgrim, 2006).

We do not aim for a "global" measure for the stability of the entire clustering algorithm (Fowlkes and Mallows, 1983; Levine and Domany, 2001) or a measure for the stability of each cluster (Monti et al., 2003), rather we want a measure for each individual clustered element – in our case: the voxel. So with inspiration from the NPAIRS framework (Strother et al., 2002, 2004), co-validation (Madani et al., 2005) and a resampling approach for cluster validation by Roth et al. (2002) we perform a half-split resampling and compare how often each voxel fall in corresponding clusters of the two split-halfs. We work from the bag-of-words and voxelization matrices, $\mathbf{X}$ and $\mathbf{Y}$, then

1. Permute randomly the rows of $\mathbf{X}$ by a permutation matrix $\mathbf{P}_r$ and take the first half of the rows of $\mathbf{X}$ for one matrix $\mathbf{X}_{1,r}$ and the second half for another matrix $\mathbf{X}_{2,r}$. Do the same operation with the same permutation matrix for $\mathbf{Y}$

$$\mathbf{X}_{1,r} = [\mathbf{P}_r \mathbf{X}]_{(1 \ldots \lceil N/2 \rceil)} \tag{8}$$
$$\mathbf{X}_{2,r} = [\mathbf{P}_r \mathbf{X}]_{(\lceil N/2 \rceil + 1 \ldots N)} \tag{9}$$
$$\mathbf{Y}_{1,r} = [\mathbf{P}_r \mathbf{Y}]_{(1 \ldots \lceil N/2 \rceil)} \tag{10}$$
$$\mathbf{Y}_{2,r} = [\mathbf{P}_r \mathbf{Y}]_{(\lceil N/2 \rceil + 1 \ldots N)} . \tag{11}$$

2. Perform nPLS on the two product matrices

$$\mathbf{W}_{1,r}\mathbf{H}_{1,r} \leftarrow \mathrm{nmf}\left(\mathbf{X}_{1,r}^{\mathsf{T}}\mathbf{Y}_{1,r}\right)$$
$$\mathbf{W}_{2,r}\mathbf{H}_{2,r} \leftarrow \mathrm{nmf}\left(\mathbf{X}_{2,r}^{\mathsf{T}}\mathbf{Y}_{2,r}\right) .$$

   We use the same number of components K as for the nPLS on the full data set.

3. Apply winner-take-all functions on the $\mathbf{H}$ matrices obtaining $\tilde{\mathbf{H}}_{1,r}$ and $\tilde{\mathbf{H}}_{2,r}$.

4. Match the rows of $\tilde{\mathbf{H}}_{1,r}$ and $\tilde{\mathbf{H}}_{2,r}$ with the Hungarian method applied on the product matrix $\tilde{\mathbf{H}}_{1,r}\tilde{\mathbf{H}}_{2,r}^{\mathsf{T}}$.

5. Record for each voxel whether it resides in the same cluster in the two independent split-halfs, i.e., whether a specific column has a non-zero value in the same rows of $\tilde{\mathbf{H}}_{1,r}$ and $\tilde{\mathbf{H}}_{2,r}$ after the rows have been matched.

6. Repeat the preceding steps R times with a new permutation matrix $\mathbf{P}_r$, $r = 1 \ldots R$, and for each voxel p count the number of times c_p it resides in the same cluster in the two split-halfs with the result in a vector $\mathbf{c}$ with length P.

The count now appears as our statistics for the stability of a voxel, and to generate a distribution for the null hypothesis for this statistics we apply a further permutation: In this permutation test we randomly permute the two half-split bag-of-word matrices, but not the voxelization matrices:

$$\tilde{\mathbf{X}}_{1,r} = \mathbf{P}_{1,r}\mathbf{X}_{1,r} \tag{12}$$
$$\tilde{\mathbf{X}}_{2,r} = \mathbf{P}_{2,r}\mathbf{X}_{2,r}. \tag{13}$$

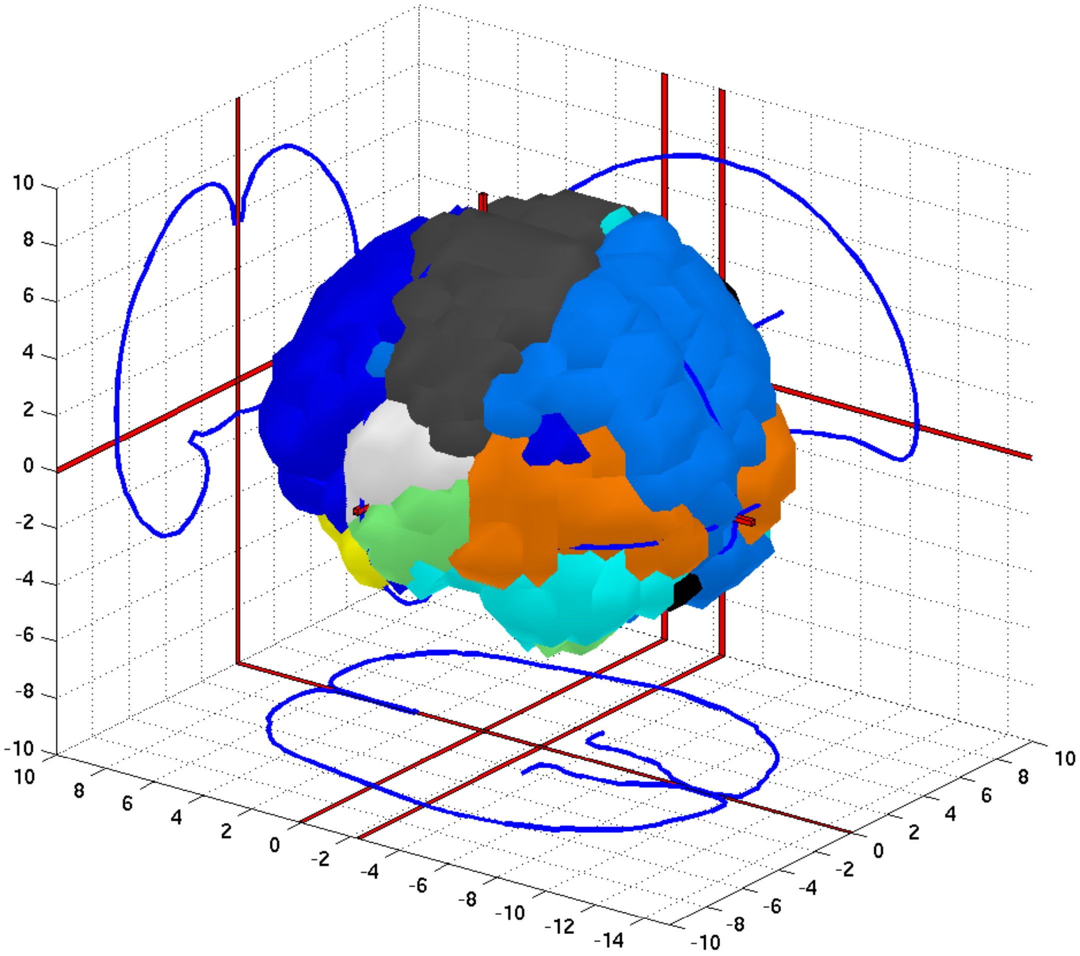

Figure 1: Surface view of functional atlas in a corner cube environment. The blue lines outline the contour of the brain from the Talairach atlas and red axes indicate the anterior and posterior commisure axes. Each colored area corresponds to a nPLS component, see Fig. 3 for a legend.

Then we perform steps 2–6 with these matrices and obtain a vector of counts $\tilde{\mathbf{c}}$. We examine the maximum in this vector and use it as the parameter in a Poisson distribution which gives us a distribution for the maximum count of a null distribution, somewhat related to the maximum permutation statistics of Holmes et al. (1996).

We repeated steps 1–5 of the cluster validation method $R = 1000$ times, each running a NMF with a maximum of 1000 iterations, i.e., a total of four million NMF iterations (2 NMFs for the half-split data and further 2 NMFs for the null distribution).

3 Results

We included data from the $N = 186$ papers contained in the Brede Database. $P = 470$ words and $Q = 2492$ voxels remained after the initial processing, and with the number of components for the nPLS set to $K = 10 \approx \sqrt{186/2}$ we obtained the results in Figs. 1, 2 and 3, where the ten components are color coded. Figs. 1 and 2 display the non-zero elements of the $\mathbf{H}$ factorization matrix (with the winner-take-all function applied) as 3-dimensional surfaces in a Talairach space corner cube environment (Rehm et al., 1998), while Fig. 3 lists the words associated with the five elements with the highest load in the $\mathbf{W}$ matrix.

Running our cluster validation method 1000 times we found 157 as the largest number of times a voxels appeared in the same cluster in both the half-split data sets of the permutation test. We use this as a parameter for a Poisson distribution $\lambda = 157$ and get 178 as a threshold value corresponding to an upper tail area in the cumulative distribution of $P \approx 0.05$. Fig. 4 displays the results after thresholding: The cluster validation retains many of the voxels (compare Figs. 1 and 4). However, some small clusters disappear, e.g., the fifth component

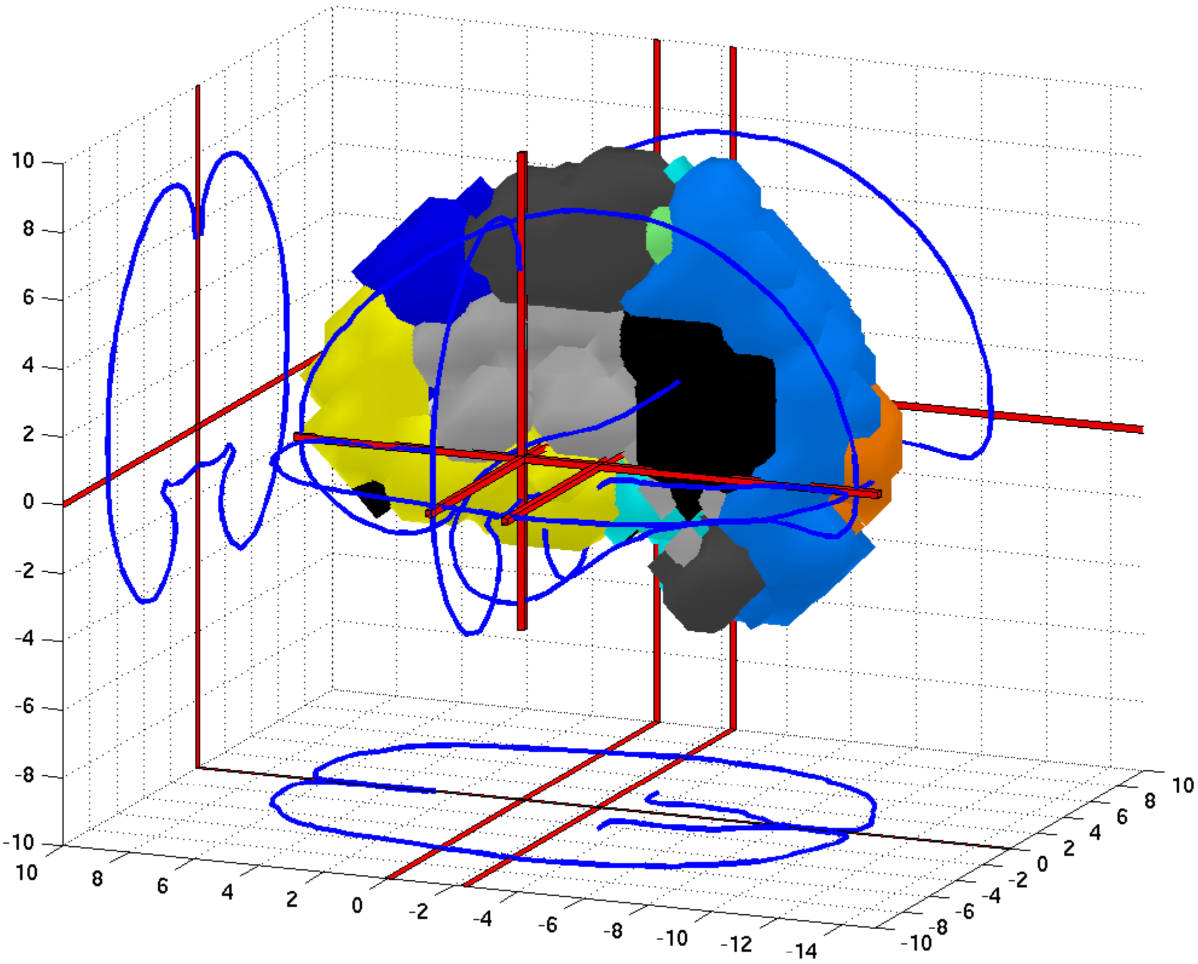

Figure 2: Functional atlas with a view towards the medial part of the right hemisphere. For the legend see Fig. 3.

(dark blue) around the left angular gyrus. Table 2 lists local maxima in the volumes of the **H** matrix, with a local maximum defined as a voxel with a load higher than its 26 neighboring voxels. The table lists only maxima that survive the cluster validation threshold. For convenience the transformed AAL atlas labels anatomically the coordinates.

The first component (black) loads predominately in the posterior cingulate area and associates with words such as "memory" and "retrieval". Table 2 reveals the highest load to be the voxel $(0, -56, 16)$. With our spatial transformation the AAL atlas labels this as in the cerebellum, but inspection of the original Talairach atlas shows this as the posterior cingulate area.

The second component (brown) clusters at the central sulcus, in the premotor cortex and the cerebellum and peaks in postcentral gyri and supplementary motor area. Sensory and motor words associates with it.

A pain topic splits between two component with the third component (grey) primarily in the anterior cingulate cortex, the thalamus and putamen (Fig. 2), and the fourth component (white) in left and right insula (Figs. 1 and 4).

The fifth component (dark blue) occupies much of the dorsal surface of the prefrontal cortex, with its peak in "area triangularis" (triangular part of the inferior frontal gyrus) and loads on words for speech/language and cognition.

Large portions of the occipital and parietal lobes incorporate a component (light blue) loaded on the words "vision", "eye" and "attention". Whereas this component falls in the "dorsal stream" area (Ungeleider and Haxby, 1994), the seventh component (cyan) relates to the "ventral stream" area with object vision words such as "faces", "recognition" and "images". Similar with the first component there is a discrepancy between the label from our transformed AAL atlas (Table 2) and visual in-

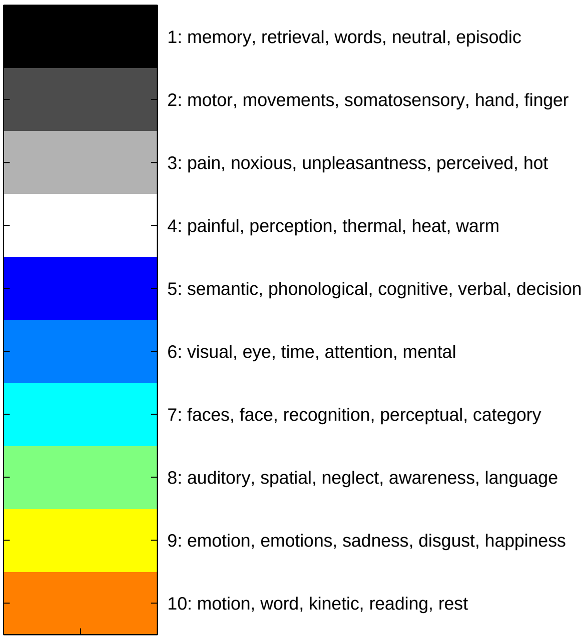

Figure 3: Ordered lists of words associated with high load on the ten components. This figure also acts as a legend for Figs. 1, 2 and 4.

spection in the orginal Talairach atlas. The latter place the maxima in the fusiform gyri.

For each of the eighth (green) and tenth (orange) components the word labeling suggests more than one brain function: The eighth component appears in the temporal gyri and labels primarily with auditory function but also incorporate spatial awareness as its other function (“spatial”, “neglect”, “awareness”). The tenth component lies posterior to the eighth component and consists of visual motion, supported by “motion” and “kinetic” high loaded words but also language supported by words “word” and “reading”.

The words of the ninth component (yellow) clearly shows an emotion function, e.g., “emotion”, “sadness” and “disgust”. It loads in mediobasal frontal cortex and goes superior in the medial area, continues posteriorly in the ventral part of the frontal cortex to amygdala towards the anterior hippocampus and to the temporal pole (Fig. 2) with peaks in the medial part of the superior frontal gyrus and the hippocampus according to our transformed AAL atlas (Table 2). The maxima falls between the lentiform nucleus, thalamus, hippocampus, closest to the amygdala and hypothalamus, when compared to the Talairach atlas with the Talairach Daemon (Lancaster et al., 2000) and the anatomical labels in the Brede Database (Nielsen and Hansen, 2002a,b).

4 Discussion

Our episodic memory retrieval component (black) in the posterior cingulate area agrees with the major review of Cabeza and Nyberg (2000) as also noted in our data mining of this particular area (Nielsen et al., 2005). Our pain components load on the areas mentioned by Ingvar (1999), though additionally we find loading in the left putamen. That pain dominates, e.g., the anterior cingulate cortex, we attribute to the many pain studies added to the Brede Database in connection with our previous studies (Nielsen et al., 2004a, 2005).

That audition loads in the temporal lobe and vision in the occipital lobe does not surprise. However, the auditory component also incorporates spatial awareness. Two spatial neglect studies in the Brede Database (Karnath et al., 2001, 2004) report Talairach coordinates in the temporal lobe at the same position as auditory studies do and the co-localization of audition and spatial awareness merely reflects this fact.

Luria (1989) stresses the major divide between a lateral and a mediobasal part of frontal cortex associated with speech activity and emotional states, respectively. We find this dichotomy with our emotion component (yellow) and the language component (dark blue). Our language component continues in the right hemisphere. Our method will label an only weakly loaded brain area provided no other brain function occupies it, so if no other brain function consistently associates with the right hemisphere it may simply be taken over by our speech and language component which also incorporates general cognition as indicated by the words “cognitive” and “decision” highly loaded on this component. Cabeza and Nyberg (2000) associate working memory with prefrontal cortex. Our database contains relatively few working memory studies. Their small number is probably responsible for the lack of an individual working memory component. Phan et al. (2002) notes that “the medial prefrontal cortex appeared to have a general role in emotion processing” and our functional atlas supports this conclusion. Steel and Lawrie (2004) consider emotion and cognition in the meta-analysis of the prefrontal cortex based on the studies identified by Cabeza and Nyberg (2000) and Phan et al. (2002), and find higher inferior medial activations for emotion induction tasks compared to cognitive tasks which our study supports. They also report difference in load in the anterior cingulate area while in our case pain occupies this areas. The limited resolution could explain why the peaks for the emotion component occur between the amygdala and hypothalamus brain structures.

	x	y	z	Load
1. memory, retrieval, words, neutral, episodic, encoding, associative, novelty				
Lobule IV, V of vermis	0	−56	16	483.07
Left fusiform gyrus	−24	−32	−8	137.83
Left gyrus rectus	0	48	−16	60.16
Right fusiform gyrus	40	−16	−24	53.68
2. motor, movements, somatosensory, hand, finger, representations, voluntary, sensory				
Left postcentral gyrus	−40	−24	48	381.54
Left supplementary motor area	0	0	56	308.61
Right postcentral gyrus	40	−24	40	117.95
3. pain, noxious, unpleasantness, perceived, hot, forearm, interaction, affective				
Left anterior cingulate gyrus	0	16	24	343.95
Left thalamus	−8	−16	8	300.22
Right thalamus	8	−16	8	288.98
Left putamen	−24	8	8	175.58
Right insula	48	0	8	158.12
4. painful, perception, thermal, heat, warm, cold, patterns, sensation				
Left insula	−40	−8	8	373.40
Right insula	40	8	8	165.82
5. semantic, phonological, cognitive, verbal, decision, fluency, syllable, pseudowords				
Left area triangularis	−48	24	16	262.73
Left superior frontal gyrus, medial part	−8	24	40	188.18
Right middle frontal gyrus, lateral part	40	32	24	137.10
Left superior frontal gyrus, medial orbital part	−8	40	0	90.17
6. visual, eye, time, attention, mental, movement, attentional, visually				
Right superior occipital	24	−64	48	218.72
Left superior parietal lobule	−24	−56	48	195.47
Left middle occipital	−24	−88	16	144.92
Right middle occipital	32	−80	16	142.20
Left calcarine sulcus	0	−80	0	132.42
Left precentral gyrus	−40	0	40	111.31
Right precentral gyrus	40	0	48	110.13
7. faces, face, recognition, perceptual, category, images, artefacts, matching				
Left Lobule VI of cerebellar hemisphere	−40	−56	−16	286.18
Right Lobule VI of cerebellar hemisphere	40	−56	−16	278.67
8. auditory, spatial, neglect, awareness, language, sounds, lesions, voice				
Right superior temporal gyrus	56	−24	8	430.81
Left middle temporal gyrus	−56	−16	0	150.40
Left middle temporal gyrus	−56	−32	8	133.13
9. emotion, emotions, sadness, disgust, happiness, emotional, facial, expressions				
Left hippocampus	−16	−8	−8	271.45
Right hippocampus	16	−8	−8	265.40
Left superior frontal gyrus, medial part	0	48	24	208.92
10. motion, word, kinetic, reading, rest, semantically, integration, sentences				
Left middle temporal gyrus	−48	−56	8	344.50
Right middle temporal gyrus	48	−64	8	239.95

Table 2: Local maxima in Talairach space in each of the 10 components. It is automatically generated and with anatomical labels from the AAL atlas.

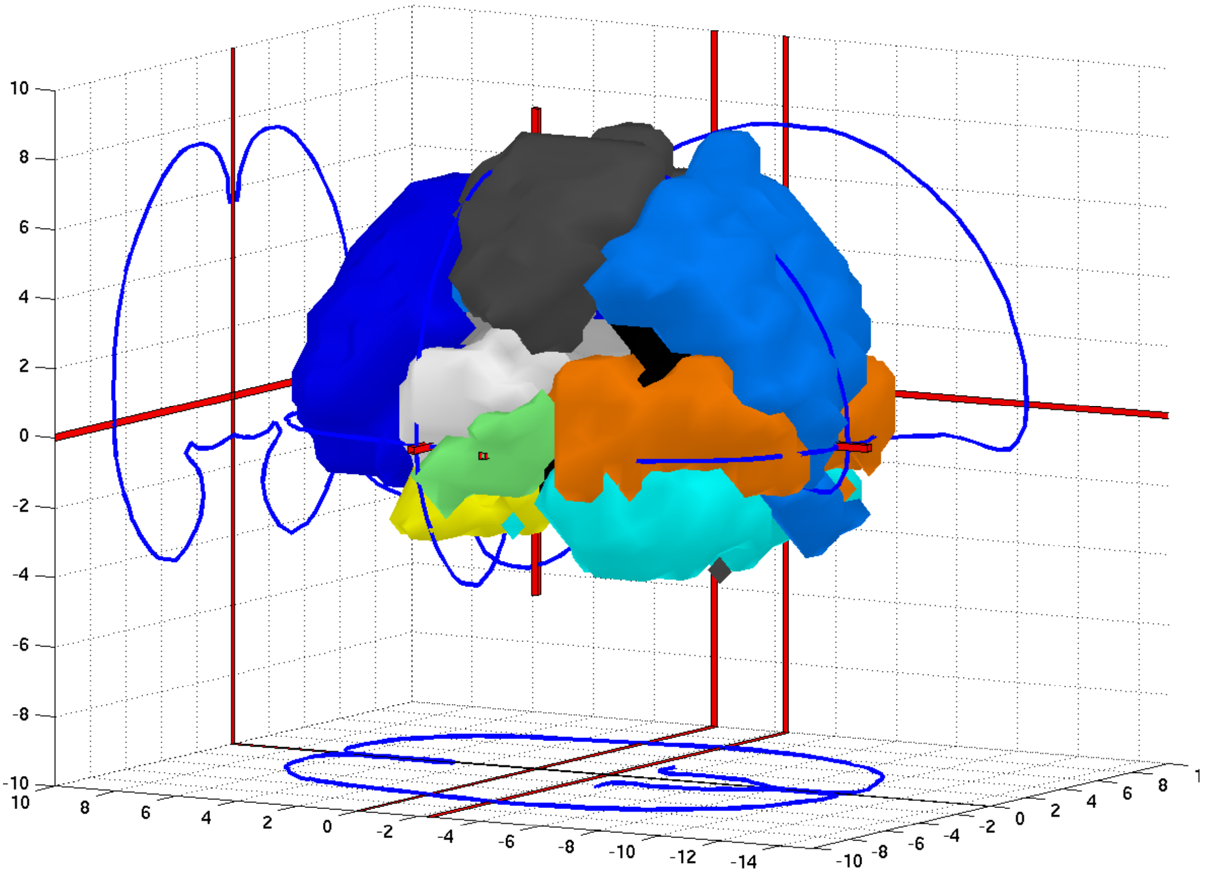

Figure 4: Functional atlas after thresholding based on cluster validation. The grey and black component in the cingulate gyrus appear in a hole on the lateral surface made from voxel that did not survive the threshold.

The bag-of-words representation of a paper makes a simplistic representation of the topic of it, e.g., words can occur in negative assertions, though this seems rarely the case. More complex representations of the abstracts could use phrase frequency rather than just word frequency or utilize natural language processing with understanding of the grammar. Furthermore, the stop word list might not be optimal, e.g., as a common personal pronoun we added the word “self” to the stop word list. However, in abstracts in the Brede database it often appears as a noun referring to a psychological concept. At present our method will thus overlook such a topic unless other words support it. On the other hand the stop word list does not contain “time” and it loads highly on the sixth component. A brief inspection shows that the word mostly finds application in contexts of experimental methodology and not so much in relation to brain function. The inclusion of this word may confound the results.

Our coarse 8 millimeter voxels and the broad 10 millimeter kernel width give a somewhat blocky and rough atlas. As the database increases in size and scope we may allow for a higher resolution and a narrower kernel. Also a new alignment between templates (Lancaster et al., 2006) may make a better match so one would expect more precise localizations. However, such detail as seen in the atlases constructed through autoradiography by Zilles et al. (2004) will probably not come about with our method. Furthermore, as more studies come into the database the eighth component might divide into an auditory and a spatial awareness component, and the orange tenth component might split into visual motion and language components, though the motion verb study of Wallentin et al. (2005) exists in the Brede Database linking these functions that one would usually regard as two entirely separate.

The threshold from cluster validation will not be valid under all interpretations of consistency. For instance, consider a voxel consistently switching between two clusters in each permutation. Among

1000 half-split iterations around 500 times the voxel will be counted as in the same cluster by our method — a number considerably higher than the null distribution threshold on 178. On the other hand, a voxel that switches uniformly and randomly between the ten components in each permutation will on average be counted 100 times as in the same cluster in the two split-halfs. The statistics of such a voxel will fall under the cluster validation threshold and be eliminated. For voxels of this type our cluster validation will make a valid cut-off.

The Brede Database has errors among the over 11,000 numbers that make up the 3-dimensional stereotaxic coordinate information. These errors arise during typing in the database or from print errors in the articles. Outlier detection can spot some of these (Nielsen and Hansen, 2002b). nPLS approximates the data and presupposes a residual regardless of whether outliers infects the data or not. Outliers would only affect the results little when the bulk of coordinates are correct.

Incorporating various kinds of studies into one model helps to balance the interpretations, e.g., the modeling of the language in Broca's area by Poldrack (2006) requires access to non-language studies to evaluate sensitivity and specificity. This type of analysis may clarify which functions make a strong involvement in an area when compared against other functions. A caveat is the limited number of studies in the Brede Database. That memory and pain tops as important components and dominate the cingulate area derives from the many studies of this kind added to the Brede Database in our previous studies. Furthermore, publication bias may make the results from neuroimaging seem more homogeneous (Phan et al., 2002). In our case it might both affect the results through under-reporting of coordinates as well as selection of words in the formulation of the abstract, e.g., when studying face recognition researchers may choose only to scan the slices of the brain containing the fusiform gyrus, and when performing an episodic memory retrieval study and finding a change in the posterior cingulate area this may be noted in the abstract simply because it falls in line with the review of Cabeza and Nyberg (2000).

The results should not be taken to support the notion of "exclusive" or "exhaustive" functional specialization (Lloyd, 2002), but rather that some areas associate particularly with a specific brain function. One area may very well be involved in several functions. Our winner-take-all function will exclusively focus on the most dominant of them and we have chosen this reductionism for the sake of conciseness and leave room for future more verbose description.

Authors disagree in how consistent functional brain imaging observations are (Lloyd, 2002; Uttal, 2002). Bearing its limitations in mind our method gives some quantitative indication of the degree of consistency in the neuroimaging literature. The components we find, however, span quite broad functions and they do not show the extent of consistency for more specialized higher-level cognitive processes, such as "moral judgment", "the anticipation of pain" and "reading body language" mentioned by Uttal (2002).

Pattern recognition should serve to make neuroinformatics databases more than just "one damn fact after another". Our functional atlas automatically analyses the content of our database and presents a view from the top. It does not rely on a taxonomy of brain functions (Nielsen, 2005) or whether it is at all possible to build a valid psychobiological taxonomy (Uttal, 2001). The bag-of-words representation makes up a description of experiments in papers on a "simple utilitarian level".

In the bioinformatics arena automated text mining has already shown its utility, e.g., with the PubGene web-service revealing association between genes (Jenssen et al., 2001), and in web-navigation automated algorithms, e.g., using PageRank (Brin and Page, 1998), play the major role compared to human generated web-indices. Our work aligns more with such services rather than the "usual" neuroimaging meta-analysis study where a human expert carefully selects the studies for investigation. One would probably trust such a meta-analysis more than our "quick and dirty" high-throughput method, but we hope that our shows utility as a data-driven overlook on functional human brain mapping.

References

Borga, M., Landelius, T., and Knutsson, H. (1992). A unified approach to PCA, PLS, MLR and CCA. Technical Report LiTH-ISY-R-1992, Computer Vision Laboratory, Department of Electrical Engineering, Linköping University, S-581 83 Linköping, Sweden. ftp://ftp.isy.liu.se/pub/bb-/TechReports/LiTHISYR1992.ps.gz.

Brett, M. (1999). The MNI brain and the Talairach atlas. http://www.mrc-cbu.cam.ac.uk/-Imaging/mnispace.html. Accessed 2003 March 17.

Brin, S. and Page, L. (1998). The anatomy of a large-scale hypertextual web search engine. *Computer Networks and ISDN*

Systems, 30(1–7):107–117. http://www-db.stanford.edu/pub/papers/google.pdf. CiteSeer: http://citeseer.ist.psu.edu-/brin98anatomy.html. Describes the Google Internet search algorithm and the PageRank algorithm.

Bro, R. and Jong, S. D. (1997). A fast non-negativity-constrained least squares algorithm. *Journal of Chemometrics*, 11(5):393–401. DOI: 10.1002/(SICI)1099-128X(199709/10)11:5¡393::AID-CEM483¿3.0.CO;2-L. http://www3.interscience.wiley.com/cgi-bin/abstract/61001584/.

Brodmann, K. (1909/1994). *Brodmann's 'Localisation in the Cerebral Cortex'*. Smith-Gordon, London, UK. ISBN 1854630288.

Cabeza, R. and Nyberg, L. (2000). Imaging cognition II: An empirical review of 275 PET and fMRI studies. *Journal of Cognitive Neuroscience*, 12(1):1–47. PMID: 10769304. http://jocn.mitpress.org/cgi-/content/abstract/12/1/1.

Carroll, J. D., De Soete, G., and Pruzansky, S. (1989). Fitting of the latent class model via iteratively reweighted least squares CANDECOMP with nonnegativity constraints. In Coppi, R. and Bolasco, S., editors, *Multiway Data Analysis*, pages 463–472. North-Holland, Amsterdam, Holland. ISBN 0444874100. Descripes a method for etimating the N-way latent class model by alternating weighted least squares with non-negativity constraints.

Chau, W. and McIntosh, A. R. (2005). The Talairach coordinate of a point in the MNI space: how to interpret it. *NeuroImage*, 25(2):408–416. PMID: 15784419.

Chein, J. M., Fissell, K., Jacobs, S., and Fiez, J. A. (2002). Functional heterogeneity within Broca's area during verbal working memory. *Physiology & Behavior*, 77(4-5):635–639. PMID: 12527011. http://www.lrdc.pitt.edu/fiez/publications-/FunctionalHeterogeneity2002.pdf.

Cutting, D. R., Karger, D. R., Penderson, J. O., and Tukey, J. W. (1992). Scatter/gather: A cluster-based approach to browsing large document collections. In Belkin, N. J., Ingwersen, P., and Pejtersen, A. M., editors, *Proceedings of the 15th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 318–329, New York, NY, USA. ACM Press. ISBN 0897915232.

Fowlkes, E. B. and Mallows, C. L. (1983). A method for comparing two hierachical clusterings (with comments). *Journal of the American Statistical Association*, 78(383):553–584.

Fox, P. T., Laird, A. R., Fox, S. P., Fox, P. M., Uecker, A. M., Crank, M., Koenig, S. F., and Lancaster, J. L. (2005). BrainMap taxonomy of experimental design: Description and evaluation. *Human Brain Mapping*, 25:185–198. DOI: 10.1002/hbm.20141.

Fox, P. T., Lancaster, J. L., Parsons, L. M., Xiong, J.-H., and Zamarripa, F. (1997). Functional volumes modeling: Theory and preliminary assessment. *Human Brain Mapping*, 5(4):306–311. http://www3.interscience.wiley.com/cgi-bin/abstract/56435/START.

Fox, P. T., Mikiten, S., Davis, G., and Lancaster, J. L. (1994). BrainMap: A database of human function brain mapping. In Thatcher, R. W., Hallett, M., Zeffiro, T., John, E. R., and Huerta, M., editors, *Functional Neuroimaging: Technical Foundations*, chapter 9, pages 95–105. Academic Press, San Diego, California. ISBN 0126858454.

Gaussier, Éric. and Goutte, C. (2005). Relation between PLSA and NMF and implications. In Baeza-Yates, R. A., Ziviani, N., Marchionini, G., Moffat, A., and Tait, J., editors, *SIGIR 2005: Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Salvador, Brazil, August 15-19, 2005*, pages 601–602. Association for Computing Machinery, New York. http://www.xrce.xerox.com/Publications-/Attachments/2005-029/gaussier05relation.pdf. ISBN 1595930345.

Goldberg, D. (1991). What every computer scientist should know about floating point arithmetic. *ACM Computing Surveys*, 23(1):5–48. DOI: 10.1145/103162.103163. ISSN 0360-0300.

Goodman, L. A. (1974). Exploratory latent structure analysis using both identifiable and unidentifiable models. *Biometrika*, 61(2):215–231.

Hofmann, T. (1999). Probabilistic latent semantic analysis. In *Proceedings of the 15th Conference on Uncertainty in AI*, pages 289–296, San Fransisco, California. Morgan Kaufmann. CiteSeer: http://citeseer.ist.psu.edu-/hofmann99probabilistic.html.

Højen-Sørensen, P. A. d. F. R., Winther, O., and Hansen, L. K. (2002). Mean field approaches to independent component analysis. *Neural Computation*, 14(4):889–918. http://isp.imm.dtu.dk/staff/winther/hojen.ica.pdf. CiteSeer: http://citeseer.ist.psu.edu/455328.html.

Holmes, A. P., Blair, R. C., Watson, J. D. G., and Ford, I. (1996). Nonparametric analysis of statistic images from functional mapping experiments. *Journal of Cerebral Blood Flow and Metabolism*, 16(1):7–22. PMID: 8530558.

Indefrey, P. and Levelt, W. J. M. (2000). The neural correlates of language production. In Gazzaniga, M. S., editor, *The New Cognitive Neurosciences*, chapter 59, pages 845–865. MIT Press, Cambridge, MA, 2nd edition. Meta-analysis of language production where the function is split in core processes, and (brain region times experiment)-matrix is built and quantitatively analyzed.

Ingvar, M. (1999). Pain and functional imaging. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 354(1387):1347–1358. PMID: 10466155.

Jenssen, T.-K., Læreid, A., Komorowski, J., and Hovig, E. (2001). A literature network of human genes for high-throughput analysis of gene expression. *Nature Genetics*, 28(1):21–28. PMID: 11326270. http://www.nature.com/ng/journal/v28/n1/full/ng0501_21.html. Gene cocitation analysis involving data from over 10 million articles from PubMed citing 13,712 human genes names obtained from HUGO, LocusLink, The Genome Database and GENATLAS. Network diagrams are automatically generated with the GraphViz software. It is the basis of the PubGene web-service.

Karnath, H.-O., Berger, M. F., Kuker, W., and Rorden, C. (2004). The anatomy of spatial neglect based on voxelwise statistical analysis: A study of 140 patients. *Cerebral Cortex*, 14(10):1164–1172. PMID: 15142954. DOI: 10.1093/cercor/bhh076. WOBIB: 133. ISSN 1047-3211.

Karnath, H.-O., Ferber, S., and Himmelbach, M. (2001). Spatial awareness is a function of the temporal not the posterior parietal lobe. *Nature*, 411(6840):950–953. PMID: 11418859. DOI: 10.1038/35082075. WOBIB: 59. ISSN 0028-0836.

Kleist, K. (1922/1934). *Kriegsverletzungen des Gehirns in ihrer Bedeutung für die Hirnlokalisation und Hirnpathologie*, volume IV, part 2 of *Handbuch der Ärztlichen Erfahrungen im Weltkriege 1914/1918*. Johann Ambrosius Barth, Leipzig, Germany.

Kötter, R., Stephan, K. E., Palomero-Gallager, N., Geyer, S., Schleicher, A., and Zilles, K. (2001). Multimodal characterisation of cortical areas by multivariate analyses of receptor binding and connectivity. *Anatomy and Embryology*, 204(4):333–349. PMID: 11720237. DOI: 10.1007/s004290100199. ISSN 0340-2061. A study on macaque brain regions using binding characteristics from 9 different ligands as well as using anatomical connectivity information. Multidimensional scaling and hierarchical clustering are used on the two receptor-times-brain-regions data matrices.

Kuhn, H. W. (1955). The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1 & 2):83–97.

Laird, A. R., McMillan, K. M., Lancaster, J. L., Kochunov, P., Turkeltaub, P. E., Pardo, J. V., and Fox, P. T. (2005). A comparison of label-based review and ALE meta-analysis in the Stroop task. *Human Brain Mapping*, 25:6–21. DOI: 10.1002/hbm.20219.

Lancaster, J. L., Tordesillas, D., Martinez, M., Mazziotta, J., Toga, A., Evans, A., Zilles, K., and Fox, P. T. (2006). Comparing Talairach and MNI reference frames using the ICBM-152 template. In Corbetta, M., Nichols, T., and Pietrini, P., editors, *NeuroImage speciel issue: Twelfth Annual Meeting of the Organization for Human Brain Mapping*, volume 31, supplement 1. Elsevier. http://www.meetingassistant.com/ohbm2006/referee/abstract_popup.php?abstractno=441.

Lancaster, J. L., Woldorff, M. G., Liotti, M., Freitas, C. S., Rainey, L., Kochunov, P. V., Nickerson, D., Mikiten, S. A., and Fox, P. T. (2000). Automated Talairach atlas labels for functional brain mapping. *Human Brain Mapping*, 10(3):120–131. PMID: 10912591. Describes the Talairach Daemon and validation of the scheme through BrainMap labels, expert annotations and an activation study.

Lange, T., Braun, M. L., Roth, V., and Buhmann, J. M. (2003). Stability-based model selection. In Becker, S., Thrun, S., and Obermayer, K., editors, *Advances in Neural Information Processing Systems 15*, pages 617–624. MIT Press, Cambridge, MA. http://books.nips.cc/papers/files/nips15/AA17.pdf.

Law, M. H. and Jain, A. K. (2003). Cluster validity by bootstrapping partitions. Technical Report MSU-CSE-03-5, Department of Computer Science and Engineering, Michigan State University, East Lansing, Michigan. http://www.cse.msu.edu/~lawhiu/papers/tr03_01.ps.Z.

Lee, D. D. and Seung, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791. PMID: 10548103.

Lee, D. D. and Seung, H. S. (2001). Algorithms for non-negative matrix factorization. In Leen, T. K., Dietterich, T. G., and Tresp, V., editors, *Advances in Neural Information Processing Systems 13: Proceedings of the 2000 Conference*, pages 556–562, Cambridge, Massachusetts. MIT Press. http://hebb.mit.edu/people/seung/papers/nmfconverge.pdf. CiteSeer: http://citeseer.ist.psu.edu/lee00algorithms.html.

Levine, E. and Domany, E. (2001). Resampling method for unsupervised estimation of cluster validity. *Neural Computation*, 13(11):2573–2593.

Lloyd, D. (2000). Terra cognita: From functional neuroimaging to the map of the mind. *Brain and Mind*, 1(1):93–116. http://ipsapp007.lwwonline.com/content/getfile/4557/1/5/abstract.htm. ISSN 1389-1987.

Lloyd, D. (2002). Studying the mind from the inside out. *Brain and Mind*, 3(2):243–259. DOI: 10.1023/A:1019984508940. A review of Uttal's New Phrenology book.

Luria, A. R. (1989). *Hjernen — Introduktion til neuropsykologien*. Nyt Nordisk Forlag Arnold Busk, Copenhagen, Denmark. ISBN 8717018978. Translated from *The Working Brain* and *Osnovy Neiropsihologii*.

Madani, O., Pennock, D. M., and Flake, G. W. (2005). Co-validation: Using model disagreement on unlabeled data to validate classification algorithms. In Saul, L. K., Weiss, Y., and Bottou, L., editors, *Advances in Neural Information Processing Systems 17*, pages 877–880, Cambridge, MA. MIT Press.

Mardia, K. V., Kent, J. T., and Bibby, J. M. (1979). *Multivariate Analysis*. Probability and Mathematical Statistics. Academic Press, London. ISBN 0124712525.

McIntosh, A. R., Bookstein, F. L., Haxby, J. V., and Grady, C. L. (1996). Spatial pattern analysis of functional brain images using Partial Least Square. *NeuroImage*, 3(3 part 1):143–157. PMID: 9345485. ftp://ftp.rotman-baycrest.on.ca/pub/Randy/PLS/pls_article.pdf.

McIntosh, A. R. and Lobaugh, N. L. (2004). Partial least squares analysis of neuroimaging data: applications and advances. *NeuroImage*, 23(1):S250–S263. DOI: 10.1016/j.neuroimage.2004.07.020.

Meilă, M. (2002). Comparing clusterings. Technical Report 418, Department of Statistics, University of Washington, Seattle, Washington. http://www.stat.washington.edu/www/research/reports/2002/tr418.ps.

Meilă, M. and Heckerman, D. (2001). An experimental comparison of model-based clustering methods. *Machine Learning*, 42(1/2):9–29.

Monti, S., Tamayo, P., Mesirov, J., and Golub, T. (2003). Consensus clustering. a resampling-based method for class discovery and visualization of gene expression microarray data. *Machine Learning*, 52(1-2):91–118. http://array.bioengr.uic.edu/~yangdai/teach/bioe594-spring03/consensusCluster.pdf.

Munkres, J. (1957). Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial and Applied Mathematics*, 5(1):32–38.

Neumann, J., Lohmann, G., Derrfuss, J., and von Cramon, D. Y. (2005). Meta-analysis of functional imaging data using replicator dynamics. *Human Brain Mapping*, 25(1):165–173. http://www3.interscience.wiley.com/cgi-bin/abstract/110474181/. ISSN 1065-9471.

Nielsen, F. Å. (2003). The Brede database: a small database for functional neuroimaging. *NeuroImage*, 19(2). http://208.164.121.55/hbm2003/abstract/abstract906.htm. Presented at the 9th International Conference on Functional Mapping of the Human Brain, June 19–22, 2003, New York, NY. Available on CD-Rom.

Nielsen, F. Å. (2005). Mass meta-analysis in Talairach space. In Saul, L. K., Weiss, Y., and Bottou, L., editors, *Advances in Neural Information Processing Systems 17*, pages 985–992, Cambridge, MA. MIT Press. http://books.nips.cc/papers/files/nips17/NIPS2004_0511.pdf.

Nielsen, F. Å. (2006). Databasing molecular neuroimaging studies. *NeuroImage*, 31:T184. DOI: 10.1016/j.neuroimage.2006.04.166. http://www2.imm.dtu.dk/pubdb/views-/edoc_download.php/4691/pdf/imm4691.pdf. Neuroreceptor Mapping 2006.

Nielsen, F. Å., Balslev, D., and Hansen, L. K. (2005). Mining the posterior cingulate: Segregation between memory and pain component. *NeuroImage*, 27(3):520–532. DOI: 10.1016/j.neuroimage.2005.04.034.

Nielsen, F. Å., Chen, A. C. N., and Hansen, L. K. (2004a). Testing for difference between two groups of functional neuroimaging experiments. In Olsen, S. I., editor, *Proceedings fra den 13. Danske Konference i Mønstergenkendelse og Billedanalyse*, number 2004/10 in DIKU Technical Reports, pages 121–129, Copenhagen, Denmark. Dansk Selskab for Automatisk Genkendelse af Mønstre, Datalogisk Institut, University of Copenhagen. http://www.diku.dk/-dsagm04/proceedings.dsagm04.pdf. ISSN 0107-8283.

Nielsen, F. Å. and Hansen, L. K. (2000a). Experiences with Matlab and VRML in functional neuroimaging visualizations. In Klasky, S. and Thorpe, S., editors, *VDE2000 - Visualization Development Environments, Workshop Proceedings, Princeton, New Jersey, USA, April 27–28, 2000*, pages 76–81, Princeton, New Jersey. Princeton Plasma Physics Laboratory. http://www.imm.dtu.dk-/pubdb/views/edoc_download.php-/1231/pdf/imm1231.pdf. CiteSeer: http://citeseer.ist.psu.edu/309470.html.

Nielsen, F. Å. and Hansen, L. K. (2000b). Functional volumes modeling using kernel density estimation. Online at http://www.imm.dtu.dk-/pubdb/views/edoc_download.php/4688/pdf-/imm4688.pdf. http://hendrix.imm.dtu.dk-/staff/fnielsen/ps/Nielsen2000Functional.ps.gz. Description of kernel density modeling of Talairach coordinates. Examples are shown with modeling the spatial distribution of Talairach coordinates from the BrainMap database based on the database labels.

Nielsen, F. Å. and Hansen, L. K. (2002a). Automatic anatomical labeling of Talairach coordinates and generation of volumes of interest via the BrainMap database. *NeuroImage*, 16(2). http://www.imm.dtu.dk/~fn/-ps/Nielsen2002Translating_abstract.ps.gz. Presented at the 8th International Conference on Functional Mapping of the Human Brain, June 2–6, 2002, Sendai, Japan. Available on CD-Rom.

Nielsen, F. Å. and Hansen, L. K. (2002b). Modeling of activation data in the BrainMap$^{\text{TM}}$ database: Detection of outliers. *Human Brain Mapping*, 15(3):146–156. DOI: 10.1002/hbm.10012. http://www3.interscience.wiley.com/cgi-bin/abstract/89013001/. CiteSeer: http://citeseer.ist.psu.edu-/nielsen02modeling.html.

Nielsen, F. Å. and Hansen, L. K. (2004). Finding related functional neuroimaging volumes. *Artificial Intelligence in Medicine*, 30(2):141–151. PMID: 14992762. http://www.imm.dtu.dk/~fn/Nielsen2002Finding/.

Nielsen, F. Å., Hansen, L. K., and Balslev, D. (2004b). Mining for associations between text and brain activation in a functional neuroimaging database. *Neuroinformatics*, 2(4):369–380. http://www2.imm.dtu.dk/~fn/ps-/Nielsen2004Mining_submitted.pdf.

Nielsen, F. Å., Hansen, L. K., and Strother, S. C. (1998). Canonical ridge analysis with ridge parameter optimization. *NeuroImage*, 7(4, part 2):S758. http://isp.imm.dtu.dk/publications-/1998/Nielsen1998Canonical.ps.gz.

Nigam, K., McCallum, A. K., Thrun, S., and Mitchell, T. (2000). Text classification from labeled and unlabeled documents using EM. *Machine Learning*, 39(2/3):1–34. DOI: 10.1023/A:1007692713085. ISSN 0885-6125.

Paatero, P. (1997). Least squares formulation of robust non-negative factor analysis. *Chemometrics and Intelligent Laboratory Systems*, 37(1):23–35. DOI: 10.1016/S0169-7439(96)00044-5.

Penrose, L. S. (1946). The elementary statistics of majority voting. *Journal of the Royal Statistical Society*, 109:53–57.

Phan, K. L., Wager, T., Taylor, S. F., and Liberzon, I. (2002). Functional neuroanatomy of emotion: A meta-analysis of emotion activation studies in PET and fMRI. *NeuroImage*, 16(2):331–348. PMID: 12030820. DOI: 10.1006/nimg.2002.1087.

Pilgrim, R. A. (2006). Munkres' assignment algorithm. Modified for rectangular matrices. Online at http://csclab.murraystate.edu/bob.pilgrim-/445/munkres.html. Accessed 2006-07-18. http://csclab.murraystate.edu/bob.pilgrim/445-/munkres.html.

Plumbley, M. D. (2003). Algorithms for nonnegative independent component analysis. *IEEE Transactions on Neural Networks*, 14(3):534–543. http://www.elec.qmul.ac.uk/people/markp/2003/Plumbley03-algorithms-c.pdf.

Poldrack, R. A. (2006). Can cognitive processes be inferred from neuroimaging data. *Trends in Cognitive Sciences*, 10(2):59–63. PMID: 16406760. http://www.poldracklab.org/pdfs/poldrack_tics_inference.pdf. Discussion of inference issues in functional brain mapping and with a use of Bayes theorem.

Rehm, K., Lakshminarayan, K., Frutiger, S. A., Schaper, K. A., Sumners, D. L., Strother, S. C., Anderson, J. R., and Rottenberg, D. A. (1998). A symbolic environment for visualizing activated foci in functional neuroimaging datasets. *Medical Image Analysis*, 2(3):215–226. PMID: 9873900. http://www.sciencedirect.com/science/article/B6W6Y-45PJY0D-7/1/48196224354fdd62ea8c5a0d85379b07.

Rigouste, L., Cappé, O., and Yvon, F. (2005). Evaluation of a probabilistic method for unsupervised text clustering. In Janssen, J. and Lenca, P., editors, *Applied Stochastic Models and Data Analysis*, pages 114–126, Brest Cedex, France. ENST Bretagne. http://asmda2005.enst-bretagne.fr/IMG/pdf/proceedings/114.pdf. ISBN 2908849151.

Roth, V., Lange, T., Braun, M., and Buhmann, J. (2002). A resampling approach to cluster validation. In Härdle, W. and Rönz, B., editors, *Proceedings in Computational Statistics: 15th Symposium Held in Berlin, Germany 2002 (COMPSTAT2002)*, pages 123–128, Heidelberg, Germany. Physica-Verlag. http://www.ml.inf.ethz.ch/publications/roth.compstat02.pdf.

Salton, G., Wong, A., and Yang, C. S. (1975). A vector space model for automatic indexing. *Communication of the ACM*, 18:613–620.

Saul, L. and Pereira, F. (1997). Aggregate and mixed-order Markov models for statistical language processing. In Zhou, J. and Church, K., editors, *Proceedings of the Fifth Workshop on Very Large Corpora*, pages 81–89, East Stroudsburg, Pennsylvania, USA. Association for Computational Linguistics. http://acl.ldc.upenn.edu/W/W97/W97-0309.pdf.

Schleicher, A., Amunts, K., Geyer, S., Morosan, P., and Zilles, K. (1999). Observer-independent method for microstructural parcellation of cerebral cortex: A quantitative approach to cytoarchitectonics. *NeuroImage*, 9(1):165–177. PMID: 9918738.

Shahnaz, F., Berry, M. W., Pauca, V. P., and Plemmons, R. J. (2006). Document clustering using nonnegative matrix factorization. *Information Processing and Management*, 42(2):373–386. DOI: 10.1016/j.ipm.2004.11.005. http://www.math.wfu.edu/publications/plemmons/documentclustering.pdf.

Shen, J. and Israël, G. W. (1989). A receptor model using a specific non-negative transformation technique for ambient aerosol. *Atmospheric Environment*, 23(10):2289–2298. DOI: 10.1016/0004-6981(89)90190-X. An early article about non-negative matrix factorization.

Steel, J. D. and Lawrie, S. M. (2004). Segregation of cognitive and emotional function in the prefrontal cortex: a stereotactic meta-analysis. *NeuroImage*, 21(3):868–875. PMID: 15006653. DOI: 10.1016/j.neuroimage.2003.09.066. A mathematical meta-analysis of Talairach coordinates in emotion and cognitive studies in the prefrontal cortex. Only one coordinate is extracted from each study. A spline-based warp transformation is applied on the coordinates so new axes correspond ot the shape of corpus callosum. A variety of tests are used on the coordinates: 2D Kolmogorov-Smirnov, MANOVA/canonical correlation analysis with weighted covariance matrix, one-dimension tests, and significance is determined by resampling. It is found that emotion coordinates tend to the inferior anterior part of the medial prefrontal cortex while cognitive tend to the posterior superior part in this region. On the lateral surface of the prefrontal cortex the emotion coordinates appear predominately in the inferior part.

Strehl, A. and Ghosh, J. (2002). Cluster ensembles — a knowledge reuse framework for combining multiple partitions. *Journal of Machine Learning Research*, 3:583–617. http://www.jmlr.org/papers/volume3/strehl02a/strehl02a.pdf. CiteSeer: http://citeseer.ist.psu.edu/557421.html.

Strother, S. C., Anderson, J., Hansen, L. K., Kjems, U., Kustra, R., Sidtis, J., Frutiger, S., Muley, S., LaConte, S., and Rottenberg, D. (2002). The quantitative evaluation of functional neuroimaging experiments: the NPAIRS data analysis framework. *NeuroImage*, 15(4):747–771. PMID: 11906218. DOI: 10.1006/nimg.2001.1034. http://linkinghub.elsevier.com/retrieve/pii/S1053811901910341. ISSN 1053-8119.

Strother, S. C., La Conte, S., Hansen, L. K., Anderson, J., Zhang, J., Pulapura, S., and Rottenberg, D. (2004). Optimizing the fMRI data-processing pipeline using prediction and reproducibility performance metrics: I. a preliminary group analysis. *NeuroImage*, 23:S196–S207. PMID: 15501090. DOI: 10.1016/j.neuroimage.2004.07.022.

Talairach, J. and Tournoux, P. (1988). *Co-planar Stereotaxic Atlas of the Human Brain.* Thieme Medical Publisher Inc, New York. ISBN 0865772932.

Tichavský, P. and Koldovský, Z. (2004). Optimal pairing of signal components separated by blind techniques. *IEEE Signal Processing Letters*, 11(2):119–122. http://itakura.kes.vslib.cz-/zbynek/pubs/SPL.pdf.

Tucker, L. R. (1958). An inter-battery method of factor analysis. *Psychometrika*, 23(2):111–136. This seems to be the first reference on the inter-battery model which is a factor analysis on two sets of variables. Remarks on this model is Gibson (1960).

Turkeltaub, P. E., Eden, G. F., Jones, K. M., and Zeffiro, T. A. (2002). Meta-analysis of the functional neuroanatomy of single-word reading: method and validation. *NeuroImage*, 16(3 part 1):765–780. PMID: 12169260. DOI: 10.1006/nimg.2002.1131. http://www.sciencedirect.com/science/-article/B6WNP-46HDMPV-N/2/-xb87ce95b60732a8f0c917e288efe59004.

Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., and Joliot, M. (2002). Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *NeuroImage*, 15(1):273–289. DOI: 10.1006/nimg.2001.0978.

Ungeleider, L. G. and Haxby, J. V. (1994). 'What' and 'where' in the human brain. *Current Opinion in Neurobiology*, 4(2):157–165. DOI: 10.1016/0959-4388(94)90066-3.

Uttal, W. R. (2001). *The New Phrenology: The Limits of Localizing Cognitive Processes in the Brain.* The MIT Press, Cambridge, Massachusetts. ISBN 0262210177.

Uttal, W. R. (2002). Response to Bechtel and Lloyd. *Brain and Mind*, 3(2):261–273.

Van Essen, D. C. (2003). Organization of visual areas in macaque and human cerebral cortex. In Chalupa, L. M. and Werner, J. S., editors, *The Visual Neurosciences*, chapter 32, pages 507–521. MIT Press, Cambridge, MA. http://brainmap.wustl.edu-/resources/VisChapter.pdf. ISBN 0262033089.

Vinod, H. D. (1976). Canonical ridge and econometrics of joint production. *Journal of Econometrics*, 4:147–166.

Wager, T. D., Phan, K. L., Liberzon, I., and Taylor, S. F. (2003). Valence, gender, and lateralization of functional brain anatomy in emotion: a meta-analysis of findings from neuroimaging. *NeuroImage*, 19(3):513–531. http://www.sciencedirect.com-/science/article/B6WNP-48J44GM-4/2-/298c8201b097f22452ec34067115ed23. Companion paper to the Phan (2002) study with kernel density estimation with a spherical uniform kernel on the stereotactic coordinates inspired by Chein (2001).

Wallentin, M., Lund, T. E., Østergaard, S., Østergaard, L., and Roepstorff, A. (2005). Motion verb sentences activate left posterior middle temporal cortex despite static context. *NeuroReport*, 16(5):649–652. PMID: 15812326. WOBIB: 175. ISSN 0959-4965.

Zilles, K., Palomero-Gallagher, N., and Schleicher, A. (2004). Transmitter receptors and functional anatomy of the cerebral cortex. *Journal of Anatomy*, 205(6):417–432. PMID: 15610391. DOI: 10.1111/j.0021-8782.2004.00357.x. http://www.blackwell-synergy.com/links/doi-/10.1111/j.0021-8782.2004.00357.x/. .